

Artificial Evolution for the Optimization of Lithographic Process Conditions

**Künstliche Evolution für die Optimierung von Lithographischen
Prozessbedingungen**

Der Technischen Fakultät der
Friedrich-Alexander-Universität Erlangen-Nürnberg

zur Erlangung des Grades

DOKTOR-INGENIEUR

vorgelegt von

Tim Fühner
aus Rheine

Als Dissertation genehmigt von
der Technischen Fakultät der
Friedrich-Alexander-Universität Erlangen-Nürnberg

Tag der mündlichen Prüfung: 24.09.2013

Vorsitzende des Promotionsorgans: Prof. Dr.-Ing. habil. Marion Merklein

Gutachter: Prof. Dr.-Ing. Dietmar Fey

PD Dr. rer. nat. Andreas Erdmann

To my late mother, Ingeborg (I miss you), my father, Herbert, and my sister, Lea.

To my late advisor, Gabriella Kókai.

Acknowledgments

I sincerely thank Professor Dietmar Fey, my first advisor, for his support and for valuable discussions that helped to improve my thesis.

I also want to extend a big thank you to Professor Gerhard Wellein for chairing the examination committee and to Professor Georg Fischer for accepting the role as an external examiner.

I would like to express my deepest gratitude to Privatdozent Andreas Erdmann, my second advisor and mentor, for his continuous support and for teaching me everything I know about lithography.

I am greatly indebted to my late advisor, Professor Gabriella Kókai. Without her encouragement and support, this thesis might not have been written. Thank you, Gabi!

I would like to thank Christoph Dürr, Stephan Popp and Sebastian Seifert for their outstanding master theses and their valuable contributions to this dissertation.

I am grateful to Dr. Jürgen Lorenz, the head of the Technology Simulation Department at Fraunhofer IISB, for his continuous support and encouragement.

I would like to thank Dr. Peter Evanschitzky for numerous discussions that greatly inspired me in my work.

Last but not least, a big thanks goes out to all my colleagues at Fraunhofer IISB for making this endeavor and my daily work an exciting and pleasant experience.

Abstract

Miniaturization is a driving force both for the performance and for cost reductions of semiconductor devices. It is therefore carried on at an enormous pace. Gordon Moore proposed and later refined an estimation stating that the minimization of costs would lead to a doubling of the density of integrated circuits every two years. And in fact, this time scale—known as Moore’s law—is still aspired at by the major players in the industry. Photolithography, one of the key process steps, has to keep up with this pace. In the past, the introduction of new technologies, including a smaller wavelength of the illumination system or higher numerical apertures (NA) of the projector, has led to a relatively straightforward downscaling approach. Today, optical lithography is confined to argon fluoride excimer lasers with a wavelength of 193 nm and an NA of 1.35. The introduction of next generation lithography approaches such as extreme ultraviolet lithography have been delayed and will not be applicable until several years from now. Further scaling hence leads to dramatically decreases process margins since patterns with dimensions of only a fraction of the wavelength have to be lithographically created.

In this work, computational methods are devised that are suited to drastically improve process conditions and hence to push resolution beyond former limitations. The lithographic process can be broadly grouped into the stepper components: the illumination system, the photomask, the projection system and the wafer stack. As shown in this dissertation, each element exhibits a number of parameters that can be subjected to optimization. To actually enhance resolution, however, a rigorous simulation and computation regime has to be established. The individual ingredients are discussed in detail in this thesis.

Accordingly, the models required to describe the lithography process are introduced and discussed. It is shown that the numerical and algorithmic implementation can be regarded as a compromise between exactness and computation time. Both are critical to obtain predictive, yet feasible approaches. Another complication is the multi-scale and multi-physics nature of the first principle process models. Although it is sometimes possible to derive individual optimization-tailored, reduced models, such an approach is often prohibitive for a concise co-optimization of multiple aspects.

In this work, we thus examine an approach that allows for a direct integration of first principle models. We investigate the use of evolutionary algorithms (EAs) for that purpose. These types of algorithms can be characterized as flexible optimization approaches that mimic evolutionary mechanisms such as selection, recombination and mutation. Many variants of related techniques exist, of which a number are considered in this dissertation. One chapter of this thesis is dedicated to the discussion of different representations and genetic operators, including motivations of the choices made for the following studies.

The lithographic process is characterized not only by a large number of parameters but can also be evaluated by a wide range of criteria, some of which may be conflicting or incommensurable—such as figures of merits like performance and manufacturability. We therefore apply a multi-objective genetic algorithm (GA) that is specifically tailored to identifying ideal compromise solutions. The characteristics of multi-objective optimization, especially when performed with evolutionary algorithms, are discussed in this thesis.

There is no such thing as a universal optimizer. EAs, for example, can be considered highly flexible, but they fail to intensively exploit local information. In an attempt to get the best of both worlds, we combine evolutionary with local search routines. We thoroughly discuss our approach to these hybrid techniques and present a number of benchmark tests that demonstrate their successful applications.

The majority of optimization problems in lithography are characterized by computationally expensive fitness evaluations. The reduction of the number of candidate solutions is therefore critical to maintain a feasible optimization procedure. To this end, we devised a function approximation approach based on an artificial neural network. Specifically, the GA population constitutes the training pattern for the

network. The resulted output is the approximated fitness function. While the global search using the GA is still conducted on the exact search space, the local search is carried out on this approximation, leading to a much reduced runtime. The efficiency and feasibility of this approach is demonstrated by a number of benchmark tests.

The algorithms, frameworks and programs developed in the scope of this work are deployed as software modules that are available through the computational lithography environment Dr.LiTHO of the Fraunhofer IISB. The general software structure is briefly discussed. In order to achieve feasible optimization runtimes, rigorous distribution and parallelization techniques need to be employed. For this dissertation, a number of different approaches are devised and discussed in this thesis.

A variety of application examples demonstrate the benefits of the devised methods. In a first set of examples, source/mask optimization problems are formulated and solved. In contrast to related work, which is mainly aimed at developing models that are specifically tailored to the underlying optimization routines, the direct approach proposed here is capable of directly employing models that are typically used in lithography simulation. A multitude of results using different problem representations is presented. Additional model options including mask topography effects are demonstrated. It is shown that the approach is not restricted to simplistic aerial image-based evaluations but is able to take the process windows and thin-film effects into account. Moreover, an extension to future resolution enhancement techniques, for example, constructively using projector aberrations, is also demonstrated.

In another example series, three-dimensional mask optimizations are performed. There, the topography including the materials of the photomask absorber are subjected to optimization. Drastically improved configurations compared to both standard optical and EUV absorbers under various illumination conditions are obtained. In order to cover all aspects of the lithography process, the last section of this thesis is devoted to the optimization of the wafer stack. As an example, the anti-reflective coating applied at the bottom of the resist to reduce standing waves in the resist profile is optimized. Different configurations including single and bi-layer coating systems are examined and optimized for, especially for double patterning applications. Significant improvements in comparison to standard stacks are shown and discussed.

The thesis finally concludes with a discussion on the different optimization strategies and the optimization and simulation infrastructure developed for this work. Advantages and challenges of the methodology are highlighted and future directions and potentials are demonstrated.

Zusammenfassung

Miniaturisierung ist sowohl für die Leistungssteigerung als auch für die Kostensenkung von Halbleiterbauelementen von großer Bedeutung und wird daher mit einer enorm hohen Geschwindigkeit betrieben. Gordon Moore leitete daraus eine Schätzung ab, die besagt, dass die Hersteller gezwungen seien, etwa alle zwei Jahre die Dichte der integrierten Schaltungen zu verdoppeln. Und tatsächlich verfolgen die Hauptakteure der Industrie dieses Ziel – bekannt als Moore's Law – noch heute. Photolithographie, einer der wichtigsten Prozessschritte, hat sich diesem Ziel unterzuordnen.

In der Vergangenheit stellte die Einführung neuer Technologiestufen, einschließlich kleinerer Wellenlängen der Beleuchtungssysteme oder höhere numerische Aperturen (NA) der Projektionsysteme, einen relativ einfachen Ansatz dar, Schaltungsstrukturen zu verkleinern. Heute allerdings muss sich die optische Lithographie auf den Einsatz von Argon-basierten Excimer-Laser mit einer Wellenlänge von 193 nm und einer NA von 1,35 beschränken. Die Einführung neuer Lithographie-Generationen, beispielsweise unter Ausnutzung extrem ultraviolett (EUV) Lichtes, verzögert sich, so dass mit ihr erst in mehreren Jahren zu rechnen ist. Eine weitere Verkleinerung der Bauelemente führt so zu einer deutlichen Verschärfung der Anforderungen an den lithographischen Prozess, da Strukturen mit Abmessungen eines Bruchteiles der zur Verfügung stehenden Wellenlänge abgebildet werden müssen.

In dieser Arbeit werden deshalb numerische Methoden entwickelt, die geeignet sind, Prozessbedingungen signifikant zu verbessern und damit Auflösungen jenseits vorheriger Beschränkungen zu erzielen. Der Lithographieprozess kann in folgende Komponenten unterteilt werden: das Beleuchtungssystem, die Photomaske, das Projektionssystem und das Schichtsystem auf der Halbleiterscheibe. Wie in dieser Dissertation gezeigt wird, weist jede dieser Komponenten eine große Anzahl optimierbarer Parameter auf. Um tatsächlich eine Verbesserung der Auflösung zu erzielen, ist jedoch der Einsatz umfassender Simulationswerkzeuge unabdingbar. Deren einzelne Bestandteile werden in dieser Arbeit erörtert.

So werden die benötigten Modelle, die den Lithographieprozess beschreiben, vorgestellt und diskutiert. Es wird gezeigt, dass die numerische und algorithmische Umsetzung aus einem Kompromiss zwischen Genauigkeit und Rechenzeit besteht. Beide Kriterien sind entscheidend bei der Entwicklung eines prädiktiven und praktikablen Ansatzes. Eine weitere Komplikation ergibt sich aus der Multi-Skalen- und Multi-Physik-Eigenschaft prädiktiver Prozessmodelle. Obwohl es bisweilen möglich ist, reduzierte Modelle für ein spezielles Optimierungsproblem zu entwickeln, eignet sich ein solches Vorgehen im Allgemeinen nicht für die gleichzeitige Optimierung mehrerer Prozessaspekte.

In dieser Arbeit wird daher ein Ansatz untersucht, der die direkte Nutzung exakter Modelle erlaubt. Als Optimierungsverfahren werden dabei evolutionäre Algorithmen (EA) entwickelt und verwendet. EAs bezeichnen probabilistische Verfahren, die evolutionäre Mechanismen wie Selektion, Rekombination und Mutation imitieren und sich durch ein hohes Maß an Flexibilität auszeichnen. Da es zahlreiche EA-Varianten gibt, widmet sich ein Kapitel dieser Arbeit der Diskussion und Untersuchung verschiedener Darstellungsoptionen und genetischer Operatoren. Dabei wird insbesondere die für diese Arbeit getroffene Auswahl motiviert.

Der lithographische Prozess umfasst nicht nur eine Vielzahl an Parametern, sondern bedarf auch der Bewertung hinsichtlich verschiedener Kriterien, von denen nicht wenige wechselseitig unvereinbar oder unvergleichbar sind. So sind beispielsweise Herstellbarkeit und Leistungsfähigkeit im Allgemeinen inkommensurabel. Daher wird ein multikriterieller genetischer Algorithmus (GA), der speziell auf die Suche nach Kompromisslösungen zugeschnitten ist, implementiert und untersucht. Die Eigenschaften von Mehrzieloptimierung, insbesondere im Zusammenhang mit evolutionären Algorithmen, werden in dieser Arbeit eingehend diskutiert.

Genau so wenig wie andere Optimierer können EAs als universell bezeichnet werden: Sie zeichnen sich zwar durch hohe Flexibilität aus, sind aber anderen Verfahren bei der intensiven Ausnutzung lokaler Informationen oft unterlegen. Eine Kombination evolutionärer und lokaler Suchalgorithmen

bietet sich deshalb an. Ein entsprechendes hybrides Verfahren wird in dieser Arbeit entwickelt, und dessen Leistungsfähigkeit wird mit Hilfe einer Reihe von Benchmark-Funktionen demonstriert.

Die Mehrzahl lithographischer Optimierungsprobleme ist durch rechenintensive Güteauswertungen charakterisiert. Die Zahl der Auswertungen muss daher auf ein Minimum reduziert werden. Es wird zu dem Zweck ein Ansatz verfolgt, bei dem die Fitnessfunktion durch eine deutlich schneller auszuwertende Ersatzfunktion genähert wird. Dabei kommt ein künstliches neuronales Netz zum Einsatz, das die durch den GA erzeugte Population aus Lösungskandidaten als Trainingsinstanzen nutzt, um so ein Modell der Fitnessfunktion zu erzeugen. Dieses Modell wird dann für eine intensive lokale Suche verwendet, während die globale GA-Suche auf der ursprünglichen, exakten Funktion durchgeführt wird. Die Effizienz und die Machbarkeit dieses Ansatzes wird an einer Reihe von Vergleichstests nachgewiesen.

Die für diese Arbeit entwickelten Algorithmen, Frameworks und Programme stehen im Rahmen der Fraunhofer IISB Lithographiesimulationsumgebung Dr.LiTHO als Software-Module zur Verfügung. Der prinzipielle Aufbau der Recheninfrastruktur wird kurz diskutiert, insbesondere im Hinblick auf die entwickelten und verwendeten Verteilungs- und Parallelisierungsverfahren, ohne die praktikable Optimierungsläufe aufgrund der hohen Rechenzeiten nicht möglich wären.

Eine Vielzahl von Anwendungsbeispielen zeigt die Vorteile der entwickelten Methoden. In einer Studie werden Beleuchtungsquellen/Photomasken-Optimierungsprobleme formuliert und gelöst. Im Gegensatz zu vergleichbaren Arbeiten, die zumeist auf vereinfachten, effizienten Modellen beruhen, wird hier ein direkter Ansatz verfolgt, der es erlaubt, exakte, in der Lithographiesimulation übliche Modelle zu verwenden. Mehrere Darstellungsvarianten werden vorgestellt und anhand zahlreicher Ergebnisse diskutiert. Die Flexibilität des Ansatzes wird unter anderem durch die Berücksichtigung von Maskentopographieeffekten demonstriert. Es wird ferner gezeigt, dass das Verfahren nicht auf die Auswertung von Luftbildern beschränkt ist, sondern auch andere Komponenten wie Prozessfenster oder Dünnfilmeffekte einbeziehen kann. Weitere Ergebnisse demonstrieren die Erweiterbarkeit des Verfahrens auf zukünftige Techniken zur Verbesserung der Auflösung, zum Beispiel, der Ausnutzung der Projektor-Aberrationskontrolle.

Ziel einer weiteren Reihe von Simulationsexperimenten ist die dreidimensionale Maskenoptimierung, in der zusätzlich zur Quellen/Masken-Optimierung auch die Topographie und die Materialeigenschaften der Photomaske optimiert werden. Dabei können deutliche Verbesserungen im Vergleich zu Standardkonfigurationen erzielt werden. Optimierungsergebnisse sowohl für optische als auch für EUV-Lithographie werden präsentiert und diskutiert. Um alle Aspekte des Lithographieprozesses abzudecken, befasst sich der letzte Abschnitt der Arbeit mit dem Schichtsystem auf der Halbleiterscheibe. Als Beispiel wird die antireflektive Beschichtung auf der Unterseite des Photolackes optimiert. Diese Beschichtung wird eingesetzt, um eine Interferenz zwischen einfallendem und rückreflektiertem Licht zu verhindern, die zu stehenden Wellen führt. Verschiedene Anordnungen, darunter Einzel- und Zweischichtsysteme, werden untersucht und verbessert. Ziel dieser Studie ist es insbesondere, die Veränderungen des Schichtsystems unter heute häufig verwendeten Mehrfachbelichtungsverfahren exakt zu beschreiben und zu verbessern.

Die Dissertation schließt mit einer Diskussion sowohl der verschiedenen Optimierungsstrategien als auch der für diese Arbeit entwickelten Optimierungs- und Simulationsinfrastruktur. Vor- und Nachteile der Methodik werden hervorgehoben und mögliche zukünftige Anwendungen und Erweiterungen vorgestellt.

Contents

Abstract	vi
Zusammenfassung	viii
Contents	ix
List of Figures	xv
List of Tables	xxi
List of Algorithms	xxiii
List of Acronyms and Abbreviations	xxv
List of Symbols	xxix
1 Introduction	1
2 Micro- and Nanolithography	5
2.1 Contact and Proximity Lithography	8
2.2 Projection Lithography	8
2.2.1 Resolution	12
2.2.2 Depth of Focus	14
2.2.3 Lithographic Process Evaluation	15
2.2.4 Illumination Properties	16
2.2.5 Mask Error Enhancement Factor	18
2.2.6 Contrast and Modulation Transfer Function	20
2.2.7 Normalized Image Log Slope	21
2.2.8 Immersion Lithography	21
2.3 Extreme-Ultraviolet Lithography	22
2.4 Electron-Beam Lithography	26
2.5 Alternative Lithography Approaches	26
2.6 Resolution Enhancement Techniques	27
2.6.1 Optical Proximity Correction	28
2.6.2 Off-Axis Illumination	32
2.6.3 Phase-Shifting Mask	38
2.6.4 Pattern Doubling	50
3 Lithography Simulation	55
3.1 Photomask Diffraction Spectrum	55
3.1.1 Finite-difference Time-Domain	56
3.1.2 Rigorous Coupled Wave Analysis and The Waveguide Method	57
3.2 Imaging	59
3.2.1 Aberrations	60
3.2.2 Vector Imaging	61
3.2.3 Jones Calculus	62

3.2.4	Stokes Parameters	63
3.2.5	Abbe's Approach	65
3.2.6	Hopkins' Approach	68
3.3	Wafer Stack Light Propagation	71
3.4	Resist Modeling	73
3.4.1	Pre-bake	74
3.4.2	Exposure Kinetics	74
3.4.3	Delay	75
3.4.4	Post-Exposure Bake (PEB)	75
3.4.5	Develop	76
4	Genetic Algorithms	79
4.1	Related Approaches	79
4.1.1	Evolution Strategy	80
4.1.2	Evolutionary Programming	82
4.1.3	Differential Evolution	83
4.1.4	Genetic Programming	85
4.1.5	Artificial Immune System	88
4.1.6	Swarm Intelligence	88
4.2	Genetic Algorithm Flow	89
4.3	Optimization Parameter Representation	90
4.3.1	Binary Versus Real-coded Genetic Algorithms	91
4.3.2	Other Representations	92
4.3.3	Schemata	93
4.3.4	Gray Code	94
4.3.5	Polyploidy	96
4.3.6	Genders	97
4.4	Constraint Handling	97
4.4.1	Penalty Methods	98
4.4.2	Feasibility-preserving Methods	99
4.4.3	Selection of an Appropriate Method	99
4.5	Selection	101
4.5.1	Stochastic Sampling With Replacement	101
4.5.2	Rank-based Selection	103
4.5.3	Stochastic Universal Sampling	103
4.5.4	Deterministic Sampling and Remainder Stochastic Sampling	104
4.5.5	Tournament Selection	104
4.5.6	Truncation Selection	106
4.5.7	Selection of an Adequate Selection Scheme	106
4.6	Recombination	107
4.6.1	n -point Crossover	107
4.6.2	Uniform Crossover	109
4.7	Mutation	109
4.8	Control Parameters	110
4.9	Diversity Preservation	112
4.9.1	Sharing	113
4.9.2	Crowding	116
4.10	Genetic Algorithms Theory	118
4.10.1	The Schema Theorem	119
4.10.2	Computational Complexity	120
4.10.3	Linkage-aware Genetic Algorithms	122
4.10.4	Exact Schema Theorems	123
4.10.5	Dynamic Models	123

4.10.6	Dynamics Systems Models	124
4.10.7	Statistical Mechanics Models	125
4.11	Multi-objective Genetic Algorithms	125
4.11.1	Multi-objective Optimization Approaches	129
4.11.2	Evolutionary Multi-objective Approaches	133
4.11.3	Non-dominated Sorting Genetic Algorithm	138
4.11.4	Strength Pareto Evolutionary Algorithm	140
4.11.5	Performance Assessment	145
5	Hybrid Local/Meta-heuristic Optimization Approach	149
5.1	Related Work on Memetic Algorithms	149
5.2	Domain of Local Search Operations	152
5.3	Memetic Algorithm Design	156
5.3.1	Trigger	157
5.3.2	Selection	159
5.3.3	Local Search Space Restrictions	160
5.4	Local Search Approaches	160
5.4.1	Sequential Quadratic Programming	161
5.4.2	Nelder-Mead Down-hill Simplex Method	167
5.5	Benchmark Problems	171
5.6	Local Function Approximation Using Artificial Neural Networks	178
5.6.1	Artificial Neural Networks	179
5.6.2	Incorporation of Function Approximation for Local Search	187
5.6.3	Benchmark Problems	189
6	Simulation and Optimization Framework	195
6.1	Lithography Simulation Environment	195
6.2	Single-objective Genetic Algorithm	195
6.2.1	Caching of Fitness Values	196
6.3	Multi-objective Genetic Algorithm Implementation	198
6.4	Problem Definition Interface	198
6.5	Task Distribution Environment	200
6.6	Brokered Network for Hybrid Optimization Approaches	202
7	Source/Mask Optimization	205
7.1	Related Work	206
7.1.1	Wavefront-based Source/Mask Optimization	207
7.1.2	Interference Map Lithography Approach	210
7.1.3	Pixelated Source Optimization	210
7.1.4	Image Fidelity Mask Optimization	213
7.1.5	Local Mask Pattern Optimization	218
7.1.6	Level-set-function-based Approach	223
7.1.7	Pixelated Chromeless Phase Lithography Mask Optimization	225
7.2	Direct Problem Statement	228
7.2.1	Parameter Representation	229
7.2.2	Objectives and Constraints	232
7.3	Results	238
7.3.1	One-dimensional Mask Representation	238
7.3.2	Two-Dimensional Mask Representation	241
7.3.3	Track/Sector Illumination Representation	246
7.3.4	Results Using the Multi-objective GA	248
7.4	Source/Mask/Projector Pupil Optimization	251
7.4.1	Problem Definition	252

7.4.2	Optimization Variables	253
7.4.3	Optimization Objectives	255
7.4.4	Results	256
8	Mask and Wafer Stack Optimization	263
8.1	Mask Absorber Optimization for Line/Space Patterns in Optical Lithography	263
8.1.1	Simulation Conditions and Optimization Variables	263
8.1.2	Optimization Criteria	265
8.1.3	Results	265
8.2	Mask Absorber Optimization for Contact Holes in Optical Lithography	269
8.2.1	Optimization Parameters	269
8.2.2	Simulation Conditions and Optimization Variables	269
8.2.3	Optimization Criteria	273
8.2.4	Results	273
8.3	Mask Absorber Optimization for EUV lithography	278
8.3.1	Optimization Variables	278
8.3.2	Optimization Criteria	280
8.3.3	Results	280
8.4	Bottom-antireflective Coating Optimization	288
8.4.1	Optimization Variables and Figure of Merit	290
8.4.2	Results	291
9	Conclusions and Future Work	301
A	Derivation of the Scalar Wave Equations	307
B	Diffraction	309
C	Scalar Imaging	313
C.1	Coherent Imaging	313
C.2	Mutual Intensity	314
C.3	Imaging Under Illumination of an Extended Incoherent Source	314
C.4	Hopkins' Formula	316
C.5	Propagation of Mutual Intensity	316
C.6	Köhler Illumination	317
C.7	Partially Coherent Imaging	317
D	Transfer Matrix Method	319
E	Parameter Control in Evolutionary Algorithms	323
E.1	Deterministic/Dynamic Techniques	323
E.2	Adaptive Techniques	323
E.3	Self-adaptive Techniques	324
F	Benchmark Functions	325
F.1	Rosenbrock's Function	325
F.2	Shekel's Foxholes Function	325
F.3	Rastrigin's Function	327
F.4	Schwefel's Function	327
F.5	Griewank's Function	329
F.6	Ackley's Function	329
	Bibliography	331

List of Own Publications	389
Index	393

List of Figures

1.1	Development of downscaling of semiconductor devices (shown as half-pitch node) and wavelength of lithographic light source over the past decades.	2
2.1	Lithographic process steps	6
2.2	Components of the lithographic projection system	9
2.3	Exposure procedure of lithographic projection tools.	11
2.4	Schematic set-up of a 248-nm projection lens system.	11
2.5	Airy disk diffraction pattern set-up.	12
2.6	Airy pattern.	13
2.7	Rayleigh criterion	13
2.8	Critical dimension measurement.	15
2.9	Lithographic process window	16
2.10	Flux density distribution of a two-dimensional grating with seven slits	18
2.11	A two-dimensional regular grating.	19
2.12	The partial coherence factor (filling factor) is given by the size ratio between source extension and projection lens.	19
2.13	Printed line width as a function of the line width on the mask.	20
2.14	Modulation transfer function for different partial coherence factors.	21
2.15	Schematic comparison of wafer stage for dry and immersion lithography.	22
2.16	EUVL projection system.	23
2.17	Laser-produced tin droplet plasma source used in EUVL.	24
2.18	EUV photomask	25
2.19	Iso/dense bias for line/space pattern. The duty ratio denotes the ratio between the feature size (line width) and the space.	28
2.20	Line shortening due to line-end rounding	29
2.21	Corner rounding at an elbow structure.	30
2.22	Use of sub-resolution assist features to maximize common process windows of isolated and dense features.	31
2.23	Comparison of different illumination settings for a lines/spaces pattern.	34
2.24	Intensity distribution under different illumination conditions.	35
2.25	Parametric representation of typical illumination types	35
2.26	Two-beam imaging.	36
2.27	Common illumination polarization states.	36
2.28	Chart of typical illumination set-ups.	37
2.29	Comparison of chromium-on-glass mask (COG) and alternating phase-shifting mask.	38
2.30	Intensity distribution under COG and AltPSM.	39
2.31	Comparison of the amplitude distribution in the spatial frequency domain of a COG and an AltPSM.	39
2.32	Process Window for COG and AltPSM.	40
2.33	Generation of phase shift in AltPSM mask. After a manufacture of a chromium-on-glass.	40
2.34	Phase conflicts with AltPSMs.	41
2.35	Double exposure for AltPSMs	42
2.36	Comparison of the diffraction order efficiencies of a COG and AttPSMs.	43
2.37	Aerial image of a 120-nm isolated bright feature, using a chromium-on-glass and an attenuated phase shifting mask with 6 percent transmission.	44

2.38	Process window for COG and AttPSM for an isolated contact hole.	45
2.39	Comparison of a chromium-on-glass and an attenuated phase-shifting mask	45
2.40	AttPSM diffraction efficiency of primary orders for different duty ratios.	47
2.41	Comparison of chromeless phase-edge PSM and phase-shutter regions to produce a line/space pattern.	49
2.42	Double patterning scheme litho-etch-litho-etch (LELE).	52
2.43	Double patterning scheme litho-freeze-litho-etch (LFLE).	53
2.44	Double patterning scheme self-aligned double patterning (SADP).	53
3.1	FDTD unit cell	56
3.2	Photomask representation of the waveguide method	58
3.3	Imaging system.	59
3.4	Polarization coordinates transform.	61
3.5	Polarization ellipse.	64
3.6	Pointcaré sphere.	64
3.7	Transmission cross-coefficient.	69
3.8	Wafer stack schematic.	72
4.1	Evolution strategy: convergence comparison of $(1 + 1)$ and $(1 + \lambda)$ selection scheme. . .	81
4.2	Genetic programming crossover operation.	86
4.3	Operation flow of genetic algorithms.	90
4.4	Example of Schemata.	94
4.5	Comparison of binary and Gray code.	95
4.6	Taxonomy of constraint handling approaches in evolutionary algorithms.	98
4.7	Schematic of different constraint handling methodologies	100
4.8	Stochastic sampling selection schemes.	102
4.9	Linear-rank-based selection.	103
4.10	Tournament selection.	104
4.11	Common crossover schemes in binary-coded genetic algorithms	108
4.12	Genetic algorithm point mutation.	110
4.13	Niche definition by distance to local optima.	114
4.14	Pareto optimality	127
4.15	Weighted sum maximizers for a bi-objective optimization problem.	132
4.16	Density estimate used in NSGA-II.	139
4.17	Strength and fitness assignment in the strength Pareto evolutionary algorithm.	142
4.18	Density approximation using a k nearest neighbor variant.	143
5.1	Local search intensity in memetic algorithms.	151
5.2	Merit function domains of genetic algorithm and local optimizer.	154
5.4	Issues of transcoding between discrete and continuous merit function domains.	155
5.5	Schematic of flow of a memetic algorithm	157
5.6	Nelder-Mead down-hill simplex operations.	170
5.7	Comparison of performances of GA and MA on two-dimensional Shekel's Foxholes function. . .	172
5.8	Comparison of performances of GA and MA on 20-dimensional Rastrigin function. . . .	173
5.9	Comparison of performances of GA and MA on ten-dimensional Griewank function. . .	174
5.10	Comparison of performances of GA and MA, using the Nelder-Mead method on 20- dimensional Griewank function.	175
5.11	Comparison of performances of GA and MA on 30-dimensional Rosenbrock function. . .	176
5.12	Comparison of performances of GA and MA on 20-dimensional Ackley function.	177
5.13	Single-layer perceptron.	180
5.14	Non-linear-separability of exclusive-or gate results.	181
5.15	Multilayer perceptron.	183
5.16	Schematic of flow of memetic algorithm with function approximation	187

5.17	Comparison of performances of approximated and non-approximated memetic algorithm on 15-dimensional Rosenbrock function.	191
5.18	Comparison of exact and neural-net-approximated Rosenbrock function graphs.	192
5.19	Comparison of performances of approximated and non-approximated memetic algorithm on 30-dimensional Rosenbrock function.	193
5.20	Comparison of performances of approximated and non-approximated memetic algorithm on 20-dimensional Rosenbrock function.	193
6.1	Binary tree representation of chromosomes.	197
6.2	OptProb result browser	199
6.3	Optimization framework.	200
6.4	Main components and call scheme of DisPyTE.	201
6.5	Twisted <i>deferreds</i>	202
7.1	Archel-based source representation.	208
7.2	Different slopes in target and actual pattern at feature edges.	212
7.3	Sigmoid function.	219
7.4	Flow of imaging and resist forward process.	219
7.5	Signed distance function.	224
7.8	Tau-map.	226
7.9	General optimization flow.	228
7.10	Mask geometry representations.	230
7.11	Illumination models.	231
7.12	Illumination representation.	231
7.13	Bit string representation of illumination parameters and rectangles.	232
7.14	Overlap healing.	233
7.15	Rectangle representation: different criteria and repair steps are applied in the “geometry stage.”	234
7.16	Mask manufacturability criterion based on morphological regularity analysis.	236
7.17	Evaluation of lithographic images.	237
7.18	Results for one-dimensional mask representation.	240
7.19	Improvement of the image performance obtained with the GA for the discrete mask representation.	241
7.20	Optimized imaging performance for 90-nm isolated lines	242
7.21	Simulated optimum main feature and scattering bar number, size and placement configuration.	243
7.22	Best optimization results achieved for a dense contact hole set-up.	246
7.23	Best optimization results achieved for a chained contact hole set-up.	247
7.24	Best optimization results achieved for an isolated contact hole.	247
7.25	Best result of a mutual optimization for both chain set-up (top) and an isolated contact hole (bottom).	248
7.26	Example of a very poorly performing solution.	249
7.27	An approximation set found during a run with 200 individuals at generation 1000.	249
7.28	A good solution manually chosen from a Pareto front approximation of the contact hole chain set-up.	250
7.29	A good solution for the isolated set-up.	250
7.30	A good solution of simultaneous optimization of chained and isolated set-up.	251
7.31	Seven line/space configurations are to be simultaneously optimized.	252
7.32	Nine constant incidence angles are assumed for the computation of the angle-specific spectra.	253
7.33	Source representation using a polar grid.	253
7.34	The scatter bar configuration for each pitch is determined by three parameters.	254
7.35	Spherical Zernike orders.	255

7.36	Sidelobe criterion.	255
7.37	Process window overlap criterion.	256
7.38	Activity diagram illustrating the simulation and evaluation flow for each candidate solution.	257
7.39	Interdependence of the source, the mask and the wavefront configuration.	260
8.1	One-dimensional bi-layer absorber stack optimization	264
8.2	Absorber material availability landscape	266
8.3	Results of dense line/space pattern absorber stack optimization.	267
8.4	Convergence of comparison of a memetic algorithm with and without approximation and genetic algorithm for absorber stack optimization.	268
8.5	Results of dense line/space pattern optimization to minimize polarization sensitivity of absorber stacks.	270
8.6	Results of dense/isolated line/space pattern co-optimization using different material modes.	271
8.7	Configuration of the non-quadratic dense and the semi-dense contact arrays.	272
8.8	Convergence graph of dense/semi-dense contact hole co-optimization.	273
8.9	Source geometry optimization variables.	279
8.10	List of materials and their corresponding refractive index (n) and extinction (k) at a wavelength of 13.5 nm.	279
8.11	Mask layer stack.	280
8.12	Process window, cross sections, and illumination settings for through-pitch optimization (AttPSM) solution with a fixed 80-nm chromium absorber.	281
8.13	Process window, cross sections, and illumination settings for through-pitch optimization (AttPSM) solution using n/k list material representation.	281
8.14	Alternative solution for through-pitch optimization (AttPSM) solution using n/k list material representation.	282
8.15	Solution candidate for through-pitch optimization (AttPSM) using n/k range representation. Refractive index is 0.9, extinction is 0.014.	283
8.16	Alternative solution candidate for through-pitch optimization (AttPSM) using n/k range representation. Refractive index is 1.1, extinction is 0.013.	283
8.17	Solution candidate for AltPSM linearity optimization.	285
8.18	Alternative solution for AltPSM linearity optimization.	286
8.19	Sensitivity of highly reflecting absorber against thickness	289
8.20	Bi-layer BARC with optimization variables.	290
8.21	Optimized single layer BARC	292
8.22	Pattern stability of optimized single-layer BARC	292
8.23	Comparison of the optimized single and bi-layer BARCs	293
8.24	CD swing curves of optimized bi-layer BARC	294
8.25	Structure of bi-layer BARC with index modulation in one period	294
8.26	Reflectivity of BARC in lithography step two	295
8.27	Bi-layer BARC optimized for the second lithography step.	297
8.28	Bi-layer BARC co-optimized for both lithography steps	298
8.29	Pareto front and examples obtained for bi-objective optimization study.	298
8.30	Different bi-layer BARC solutions obtained with the multi-objective optimization approach.	300
B.1	Schematic of initial point leading Helmholtz-Kirchhoff integral theorem.	309
B.2	Fresnel-Kirchhoff diffraction principle.	310
C.1	Intensity distribution in dependence of mutual coherence of two points.	314
C.2	Mutual intensity of points remote from extended incoherent source.	315
C.3	Propagated mutual intensity.	316
D.1	Thin-film layer consisting of one film, the medium of incidence and the substrate.	319

F.1	Rosenbrock's function	326
F.2	Variant of Shekel's foxholes function	327
F.3	Rastrigin's function	328
F.4	Schwefel's function	328
F.5	Griewank's function	329
F.8	Ackley's function	330

List of Tables

2.1	History of illumination sources in projection lithography.	10
4.1	Domination relations on individual objective vectors and corresponding Pareto set approximation interpretations (reproduced from Zitzler et al., 2003).	145
5.1	Comparison of incremental learning results for different network topologies on Rosenbrock's function.	189
5.2	Comparison of learning success for different algorithms and parameters on Rosenbrock's function.	190
7.1	GA settings for both the continuous and discrete mask representation.	239
7.2	Illumination source layouts and corresponding parameters.	242
7.3	Results for different illumination settings for a chain set-up.	244
7.4	Result of the mask/source optimization for dense, chain, and isolated contact holes. . .	245
7.5	Evolution of best candidate solutions for a number of example generations.	257
7.6	Example solutions of the multi-objective genetic algorithm.	261
8.1	Dense line/space AttPSM stack optimization using a memetic algorithm with function approximation.	269
8.2	Evolution of best candidate solutions for a number of example generations of the co-optimization of dense/semi-dense contact holes.	274
8.3	Comparison of the best individuals of dense/semi-dense contact hole co-optimizations using different material modes.	277
8.4	Results for lines/spaces linearity optimization of AttPSM.	284
8.5	AttPSM contact hole optimization results.	287

List of Algorithms

3.1	Schematic Abbe vector imaging approach.	67
3.2	Resist image computation.	73
4.1	Basic Differential Evolution Algorithm	84
4.2	Tournament Selection	105
4.3	Restricted tournament selection	118
4.4	Basic non-dominated sorting algorithm	135
4.5	NSGA population ranking	138
4.6	NSGA population reduction	140
4.7	SPEA Fitness Assignment	141
4.8	SPEA Archive Reduction	144
5.1	Basic SQP Algorithm	162
5.2	Nelder-Mead Down-hill Simplex	168
5.3	Backpropagation Learning for ANN with one hidden layer	184
5.4	Resilient backpropagation algorithm (Rprop)	186
6.1	Fitness Cache	197
7.1	Multi-stage Optimization Strategy	207
7.2	Local variation optimization scheme.	216

List of Acronyms and Abbreviations

AF	assist feature
AltPSM	alternating phase-shifting mask
ANN	artificial neural network
AttPSM	attenuated phase-shifting mask
BARC	bottom anti-reflective coating
BCGA	binary-coded genetic algorithm
BFGS	Broyden-Fletcher-Goldfarb-Shanno
BIM	binary mask
CAR	chemically amplified resist
CD	critical dimension
CDF	cummulative distribution function
CEL	contrast enhancement layer
COG	chrome-on-glass
CPL	chromeless phase lithography
DAG	directed acyclic graph
DE	differential evolution
DisPyTE	a distributing Python tasks environment
DM	decision maker
DOE	diffractive optical element
DOF	depth of focus
DRM	development rate monitor
DUV	deep-ultraviolet
EA	evolutionary algorithm
EDA	electronic design automation
EL	exposure latitude
EMF	electro-magnetic field
EP	evolution strategy
ES	evolution strategy
EUV	extreme-ultraviolet
EUVL	extreme-ultraviolet lithography
FDTD	finite-difference time-domain
FEM	finite element method
FFT	fast Fourier transform
FIB	focused ion beam
FIT	finite integration technique
GA	genetic algorithm
GiNGA	GiNGA is a niching genetic algorithm
GP	genetic programming

HTTP	hypertext transfer protocol
HVM	high-volume manufacture
IC	integrated circuit
ILS	image log slope
ILT	inverse lithography technique
IML	interference mapping lithography
IPL	ion beam projection lithography
IPS	intensity in the preferred polarization state
ITRS	international technology roadmap for semiconductors
KKT	Karush-Kuhn-Tucker
LELE	litho-etch-litho-etch
LER	line-edge roughness
LFLE	litho-freeze-litho-etch
LICQ	linear independence constraint qualification
LisBON	LisBON is brokered optimization network
LLE	litho-litho-etch
LO	local optimization
LP	linear program
LPP	laser-produced plasma
LS	local search
LTEM	low thermal expansion material
LVM	low-volume manufacture
LWR	linewidth roughness
MA	memetic algorithm
MEEF	mask error enhancement factor
MEF	mask error factor
MLP	multilayer perceptron
MOEA	multi-objective evolutionary algorithm
MOGA	multi-objective genetic algorithm
MoSi	molybdenum-silicon
MPI	message passing interface
MSE	mean-squared-error
MTF	modulation transfer function
NA	numerical aperture
Nd:YAG	neodymium-doped yttrium aluminium garnet
NEMS	nano-electromechanical systems
NFL	no free lunch
NGL	next-generation lithography
NILS	normalized image log slope
NLP	non-linear problem
NM	Nelder-Mead
NPGA	niched-Pareto genetic algorithm
NSGA	non-dominated sorting genetic algorithm
OPC	optical proximity correction
OR	operations research
PAC	photo-active compound

PAES	Pareto-archived evolution strategy
PAG	photo-acid generator
PDE	partial differential equation
PESA	Pareto envelope-based selection algorithm
PML	perfectly matched layer
PSF	point spread function
PSM	phase-shifting mask
PSO	particle swarm optimizer
PW	process window
Pythmea	Python multi-objective evolutionary algorithm
QP	quadratic program
RCEL	reversible contrast enhancement layer
RCGA	real-coded genetic algorithm
RCWA	rigorous coupled wave analysis
RET	resolution enhancement techniques
RMSE	root mean square error
RPC	remote procedure call
Rprop	resilient backpropagation
RTS	restricted tournament selection
SA	simulated annealing
SB	scattering bar
SEM	scanning electron microscope
SMO	source/mask optimization
SMPO	source/mask/projector pupil optimization
SOCS	sum of coherent systems
SPEA	strength Pareto evolutionary algorithm
SQP	sequential quadratic program
SRAF	sub-resolution assist features
SVD	singular value decomposition
TARC	top anti-reflective coating
TCC	transmission cross-coefficient
TE	transverse electric
TM	transverse magnetic
TMM	transfer matrix method
TS	tabu search
TSP	traveling salesman problem
UV	ultraviolet
VEGA	vector-evaluated genetic algorithm
WLOG	without loss of generality
XML	extensible mark-up language
XOR	exclusive-or

List of Symbols

A	Amplitude of wave.
$\arg \max$	Maximizer. $\mathbf{x} = \arg \max f$: \mathbf{x} that maximizes f .
$\arg \min$	Minimizer. $\mathbf{x} = \arg \min f$: \mathbf{x} that minimizes f .
$*$	Conjugate complex. $z = a + bi$, $z^* = a - bi$.
$f * g$	Convolution. Bilinear operator. For example, $f, g : \mathbb{C} \rightarrow \mathbb{C}$, $(f * g)(t) := \int f(\tau)g(t - \tau)d\tau = \int f(t - \tau)g(\tau)d\tau$
\mathbb{F}^n	n -dimensional coordinate space, where \mathbb{F} is the field such as \mathbb{R} or \mathbb{C} .
$f \star g$	Cross-correlation. Bilinear operator. For example, $f, g : \mathbb{C} \rightarrow \mathbb{C}$, $(f \star g)(t) := \int f(\tau)g(t + \tau)d\tau = \int f(t + \tau)g(\tau)d\tau$
$\nabla \times \mathbf{F}$	Curl of three-dimensional vector field \mathbf{F} .
$:=$	$a := b$: a is defined as b .
$\nabla \cdot \mathbf{F}$	Divergence of three-dimensional vector field \mathbf{F} .
\succ	Dominance relation. $\mathbf{x} \succ \mathbf{y}$: \mathbf{x} dominates \mathbf{y} .
\succ_ϵ	Epsilon dominance relation. $\mathbf{x} \succ_\epsilon \mathbf{y}$: \mathbf{x} epsilon-dominates \mathbf{y} .
\succeq	Weak dominance relation. $\mathbf{x} \succeq \mathbf{y}$: \mathbf{x} weakly dominates \mathbf{y} .
$\mathcal{F}(f)$	Fourier transform of function f .
∇	Gradient of scalar function. Jacobian matrix of a vector function; i.e., given $\mathbf{f} : \mathbb{R}^n \rightarrow \mathbb{R}^m$, $\nabla \mathbf{f}(\mathbf{x}) := [\nabla f_1(\mathbf{x}), \dots, \nabla f_m(\mathbf{x})]^T$
\circ	Hadamard or Schur product. $\mathbf{x} \circ \mathbf{y}$: component-wise product of two vectors.
\Longleftrightarrow	$A \Longleftrightarrow B$. A if and only if B .
\Longrightarrow	Implication. $A \Longrightarrow B$: A implies B .
\parallel	Incommensurability. $\mathbf{x} \parallel \mathbf{y}$: neither \mathbf{x} dominates \mathbf{y} nor vice versa.
\mapsto	Map. $x \mapsto x^2$: x maps to x^2 .
\rightarrow	$f : \mathbb{R}^2 \rightarrow \mathbb{R}$: f is a function on \mathbb{R}^2 to \mathbb{R} .
\mathbf{A}	Matrix. $A = [a_{i,j}]_{m \times n}$. A is an m (rows) \times n (columns) matrix.
\mathcal{N}	Random number drawn from a normal distribution.
\triangleright	Outperformance with respect to dominance. $\mathfrak{X} \triangleright \mathfrak{Y}$: optimum set \mathfrak{X} outperforms \mathfrak{Y} .
μ	Pay-off, merit. $\mu(\mathbf{x})$ merit due to variable values \mathbf{x}
ϕ	Phase of wave.
p	Pitch between features (e.g., on the photomask).

\hat{p}	System-independent pitch between features (e.g., on the photomask). $\hat{p} := p \text{NA}/\lambda$.
P	Probability. $P(A)$: probability of event A.
\mathcal{R}	Range.
\mathbb{R}^-	Set of negative real numbers.
\mathbb{R}^+	Set of positive real numbers.
\mathbb{R}_0^+	Set of non-negative real numbers.
s.t.	Subject to (constraints).
\mathcal{T}	Transmission cross-coefficient. Could be described as the modulation transfer function of partial coherent imaging systems
T	Transposed vector/matrix. \mathbf{x}^T : transposed of \mathbf{x} .
\mathcal{U}	Random number drawn from a uniform distribution.
λ	Wavelength. Typically given in micrometers or nanometers.

1 Introduction

The semiconductor industry is among the most competitive sectors. This fact becomes most evident when considering the significant changes of the corporate landscape over the past decades, seeing a large number of companies either abandoning the integrated circuit (IC) business or even defaulting. But at the same time, there are massive concentrations with players successfully increasing their market share. One example for this development is also the phenomenon of companies going “fab-less,” enterprises purely devoting their business to the design and marketing of IC products. Started in the mid-1980s, this trend continues and is paralleled by an emergence of large foundries that completely focus on manufacturing.

The demand for semiconductor products, such as microprocessors or memory, is steadily increasing, offering distinguished market opportunities. The participation in this market, however, is inhibited by high technological prerequisites, associated with both the requirement of a large investment and an elevated risk of failure. On the other hand, given a suited chip design and a well-adjusted fabrication process, an extremely high-volume manufacturing can be established, providing a high potential not only for revenue but also for claiming a significant share of the market. Since all participants follow this strategy, even the greatest demand is eventually saturated.

IC manufacturers are hence even more anxious to continuously improve their fabrication process, to increase their market share by offering a superior or cheaper product or, negatively put, not to fall behind their competitors. The semiconductor sector qualifies as cut-throat competition, after all. This is fueled by the specific quality of ICs for their customers, which in general is completely absorbed in quantitative measures such as speed (i.e., operations per time unit), storage capacity or power efficiency.¹ From the perspective of design and production, these properties directly translate into two parameters: a progressive down-scaling of circuits and devices and an increase of devices (transistors) per circuit, to be achieved, of course, at the lowest possible cost per piece.

These goals are connected with both enormous investments and a huge technological effort. And their achievement entails the dependence on many parties including research and development facilities and component and equipment suppliers, many of which follow their own agenda, specifically, their own profit interests. Concerting all interdependent participants is hence a tremendous challenge, nevertheless inevitable in order for the endeavor to be substantially successful (which still does not necessarily turn it into economical success). The popularity of Moore’s law is a vivid example for the demand for such a plan. In 1965, Moore, retrospectively assessing the past years of development, noted that the number of transistors had doubled every two years and speculated that this development would continue for at least another decade (Moore, 1965). This statement was of course derived from the achievements of successful IC companies. In other words, it was the marker for success at that time. The attempt to extrapolate this success ratio was instantaneously reinterpreted as the pace to which a manufacturer has to accommodate in order to be commercially successful. Consequently, nonetheless not lacking a certain absurdity, this rate is still understood as a law that companies have to comply with in order not to perish. In that sense, it also serves as a planning instrument by which different players mutually impose their commitment on each other. The need to comprehensively plan and synchronize research and development activities is further documented by the international technology roadmap for semiconductors (ITRS), an experts-devised collection of recommendations for the achievement of anticipated future technological options, called technology nodes. The proposals of this roadmap target all aspects of semiconductor design and fabrication including lithography, the pattern duplication process step. Lithography can be considered one of the most critical steps since its resolution limit also defines the minimum feature sizes that

¹Taking actual quality measures such as reliability or durability for granted, which in fact, given the typical short life cycles or, economically, periods of amortization, even play a secondary role. Of course, there are also exceptions, which, however, apply mostly to niche products (e.g., real-time-capable processing units or controllers for non-standard tasks).

can be achieved in subsequent process steps. It is thus apparent that a considerable effort is made for lithography to meet the resolution requirements.

Enabling a high throughput, the lithographic pattern transfer is currently performed optically. We shall see in the next chapter that resolution of optical lithography is determined by a number of factors, such as the wavelength of the illuminator, characteristics of the projection system and the properties of the photosensitive material exposed by the lithographic system. In the past, meeting the size criteria was relatively simple since, for example, the employed wavelength was significantly smaller than the structures to be patterned. As shown in Figure 1.1, this situation has drastically changed over the past years. Now, lithography operates in a regime of target features significantly below the wavelength. This circumstance severely narrows the lithographic process margin. Beginning in the mid 1990s, photolithographers thus had to start to incorporate resolution enhancement techniques in order to guarantee for an acceptable image performance. Due to these measures, today the photomask (the pattern template) drastically differs from the design intent. Similarly, sophisticated developments had to be taken on the wafer side, leading, for example, to numerous innovations related to photo-chemical materials. As the scaling continues, photolithography reaches its ultimate physical limits, and the introduction of a successor is urgently needed but delayed due to technological issues. For example, one of the most promising candidates to replace the current exposure tools—using an Argon fluoride laser with a wavelength of 193 nm—is a jump to a wavelength of about 13.5 nm, termed extreme-ultraviolet (EUV). Under those conditions, however, refractive systems are infeasible, due mainly to a lack of suited materials. The catoptrics systems developed today are principally operable. A number of severe technological problems are still under investigation and their solution is still unclear.




Figure 1.1: Development of downscaling of semiconductor devices (shown as half-pitch node) and wavelength of lithographic light source over the past decades.

Originally intended as interim solutions, techniques in which the pattern transfer is decomposed into several steps and thus circumvent the resolution limit are currently being used and further developed. This, however, degrades the throughput capabilities of lithography and also drastically increases the cost of ownership of the lithographic process equipment, undermining the original goal of a low cost per product unit.

It goes without saying that IC manufacturing depends on a massive adoption of advanced computational techniques and machinery. This is of course also true for lithography. Early lithography

simulation programs were available as early as in the late 1970s² and turned into universal tools—not only for research and development but also with in-process applications—by the early to mid-1990s. Today, lithography simulation can be considered highly mature. The employed models are both accurate and efficient. Enormous advances in computing power due to the progress the semiconductor industry brought forward itself but also to algorithmic advances, for example, in exploiting parallel architectures or specialized hardware such as graphics processors, allow for an application of rigorous models also on larger scales. On the other hand, paying tribute to the inevitable limitations posed by the narrowing process margins, lithographic effects are taken into account by all major electronic design automation (EDA) software vendors, which provide optical proximity correction routines or even functions turning a layout into a “litho-friendly” one based on approximate lithographic models. But even that approach, already highly computationally challenging, is facing its limits because it does not account for novel resolution enhancement options and because the simplified assumptions are not valid in a regime close to the physical limits. The demands toward current lithography led to an increasing awareness of this fact over the passed years. A number of attempts to stringently optimize aspects of the lithographic process have consequently emerged.

In this thesis, we aim at identifying optimization problems and propose their direct solution. With the interaction of multiple influencing factors, all of which generally involve advanced numerical models, chances to gather a priori knowledge of the search space associated with such optimization problems diminish. Moreover, a number of properties, including multimodality, multiple, incommensurable objectives and non-linear constraints pose significant problems for their solution. Over the past decades a class of global optimization routines, termed evolutionary algorithms, have been successfully applied to such problems. Here, we have therefore explored their application to various lithographic process optimization problems.

The remainder of this thesis is organized as follows: The next chapter briefly introduces the state-of-the-art of lithography, specifically reviewing critical issues and their countermeasures such as resolution enhancement techniques, optical proximity correction and pattern multiplication. These techniques can be considered the toolbox of our optimization effort. Chapter 3 provides an introduction into the numerical models that are applied in this work to simulate the lithographic process. In Chapter 4, evolutionary and especially genetic algorithms are discussed, emphasizing some of their critical aspects including decision space representations, constraint handling techniques and multimodal search spaces. The last section of this chapter is devoted to the emerging field of multi-objective approaches. It comprises a brief introduction into the theory of multi-objective optimization and a comprehensive survey of its realization using evolutionary algorithms, which serves as a motivation for the selection of the specific approaches used in this work. Chapter 5 introduces our approach to hybridizing evolutionary algorithms and local search approaches. After a brief summary of related work, our design considerations are explained, and the resulting procedure is verified on a number of benchmark problems. In the last part of the chapter, an extension to the hybrid algorithm, using a neural-network-based function approximation technique is discussed, and benchmark results are shown. The software components involved in this work are demonstrated in Chapter 6. As a first example, the co-optimization of lithographic sources and masks is presented in Chapter 7. Several alternative problem formulations are proposed and compared. A different perspective on the lithographic process is taken for the optimization procedures demonstrated in Chapter 8. There, both photomasks and the wafer stack are rigorously optimized for ideal process conditions. The thesis concludes with a brief summary of the findings and an outlook on potential future directions (Chapter 9).

²Early photoresist models were even put forth by the mid-1970s by seminal work of Dill and co-workers (see Section 3.4).

2 Micro- and Nanolithography

The umbrella term lithography describes pattern transfer techniques, often used for mass reproduction purposes. The term is derived from the Greek words λίθος (lithos), stone and γράφειν (graphein), to write. It originally refers to a printing method using a limestone or metal plate, which was invented in 1798 by Johann Alois Senefelder, a German playwright, and which is still widely used, for example, for high-volume print products in the form of offset lithography.

With the invention of ICs in the late 1950s, lithography (or more precisely, microlithography) was also established as the term describing the patterning techniques used in IC manufacturing. Depending on the application—especially on the size of the target patterns and the quantity of the manufactured components—, different lithography process options can be employed. The predominant approach applied in the semi-conductor fabrication process is photolithography or optical lithography. A template, called *photomask*, is optically imaged into a photosensitive layer on the wafer, called *photoresist*. After develop, the so obtained pattern on the wafer can be used as a mask for subsequent process steps, such as etching or deposition.

The individual steps involved in the lithographic pattern transfer are schematically depicted in Figure 2.1 and can be summarized as follows (cf. Levinson, 2001; Mack, 2007; Erdmann, 2010):

- (a) **Cleaning and Surface Preparation** In a first step, the wafer must be prepared in order for the photosensitive substance, the photoresist, to adhere to it. For that purpose, the wafer is dehydrated in a heating step. Then an adhesion agent is applied. Very commonly, hexamethyldisilazane (HMDS) is used. The result is a hydrophobic film on the wafer.
- (b) **Resist Coating** After application of the adhesion layer, the resist is spun onto the wafer: The liquid resist is dispensed on the wafer after which the wafer is rotated a several thousand revolutions per minute. Depending on the viscosity of the resist and on the spin speed, a more or less thin resist film remains. The designated thickness of the resist strongly depends on the process and the circuit layer for which it is applied. Critical layers, for example, require a thickness of some ten nanometers.
- (c) **Pre-bake and Alignment** In a next step, the density of the resist is increased by a pre- or soft-bake step. An additional effect of this process step is the reduction of residual solvent, thus improving the adhesion of the resist and decreasing its sensitivity against contamination.

Since the manufacturing of modern integrated circuits consists of a number of different patterning steps, most of which start with lithography, the wafer has to be re-aligned in order to guarantee for a precise overlay of all layers.

- (d) **Exposure** The previous steps prepare what can be said to be the most critical stage, the exposure step. Here, the photochemical property of the photoresist, previously coated onto the wafer, is exploited. A light pattern in the photoresist is created through an optical system (different options are discussed later), either changing the solubility of the exposed spots of the photoresist or, in chemically amplified resists, deprotecting the acid that acts as a catalyst during the subsequent bake step.

Photoresists can be classified according to their behavior upon exposure: (1) Unexposed areas of a *positive tone resist* exhibit a significant lower solubility than exposed areas in the presence of developer. (2) A *negative tone resist* shows the exact opposite behavior: While initially being highly soluble by the developer, it shows a drastically reduced solubility at exposed spots. Negative resists were commonly used in the era of contact and proximity lithography (see Section 2.1). However, higher resolution requirements and a number of disadvantages of the negative resist process—including swelling because of solvent absorption by the resist, a high sensitivity to oxygen, and the




Figure 2.1: Lithographic process steps

need for toxic organic solvents—led to a changeover to positive tone resists in the late 1970s: For many years, a photoresist consisting of a phenolic novolak resin, a diazonaphthoquinone (DNQ) type dissolution inhibitor, and an organic casting solvent were chiefly used in the emerging projection lithography processes, for which mercury g- and i-line source were utilized (see Section 2.2). Although the positive tone process has no principal advantage, the novolak-DNQ resists show superior properties both in terms of handling and resolution compared to the formerly used negative tone ones.

With the introduction of deep-ultraviolet (DUV) lithography, which initially meant a significant reduction of the exploitable dose because of weaker illumination sources, also a new resist type was required: *Chemically amplified resists (CARs)* contain photo-acid generators (PAGs) which produce acid upon exposure, which in turn cleaves protection groups from the base polymer. This reaction, which is often triggered by a heating process, as explained in the next paragraph, is catalytic—the acid is not consumed. Thus the same acid deprotects several polymer molecules. Compared to a non-chemically-amplified resist process, lower exposure doses can be applied to attain photochemical of the same magnitude in the resist. Positive tone chemically amplified resists still prevail in state-of-the-art projection lithography processes that employ excimer laser based DUV sources.

- (e) **Post-Exposure Bake (PEB)** In most modern lithography processes, exposure is followed by a baking step, whose function strongly depends on the resist material. When applied to negative tone resists, the PEB step activates the cross-linking agent to reduce the solubility of the exposed areas of the resist. In the case of chemically amplified resists, the deblocking reaction is activated by the post-exposure bake. The photo-acid generated upon exposure does not itself change the solubility of the resist. Instead it serves as a catalyst by attacking the protection groups of the polymer.

Another important feature of PEB is its ability to counteract interference-induced artifacts, called standing waves, through diffusion. Standing waves are caused by the fact that light traveling through the resist may be reflected from any interface below, and hence interfere with the incoming light. The resulting interference patterns are photochemically materialized in the resist and may drastically degrade the quality of the target pattern. To reduce standing waves, the substrate is often coated with an additional antireflective layer, the bottom anti-reflective coating (BARC), before the resist is spun on. Although reducing the standing waves effect significantly, this coating does not eliminate it entirely. Diffusion, for example of the photo-acid in the case of chemically amplified resists, may be intentionally evoked by means of PEB to blur standing waves and to so reduce their impact on succeeding process steps.

Both the ideal baking temperature and its duration strongly depend on the used resist material and on the lithographic process. In state-of-the-art processes, hot plate heating with a temperature of about 375–400 K and a baking time of about 100 s are typically used.

- (f) **Develop** During develop a solvent is applied onto the resist to remove unexposed resist areas in the case of a negative tone process or exposed areas of a positive tone resist. The paramount development technique used in modern clustered single-wafer processes is called puddle development. There, developer is sparsely applied onto the slowly spinning wafer. The wafer is then stopped for a defined development time, after which the solvent and with it the solute resist is rinsed off, leaving only the unexposed profile. In the case of a negative tone resist, the behavior is reversed. That is, the developer attacks the unexposed photoresist.
- (g) **Hard Bake** The hard bake (sometimes also post-bake) step is applied to fortify the resist profile after develop for subsequent process steps such as etch or implant. Depending on the resist type, high temperatures of up to 475 K are used, leading to a cross-linking of the polymers and an evaporation of its organic materials. In addition to a solidification and an increase of the thermal stability of the resist profile, this is also a measure to prevent contamination in the following process steps, many of which are carried out in a vacuum environment.

- (h) **Process** The so prepared resist profile can be considered a mask for the following pattern transfer steps such as etch, deposition or implant.
- (i) **Strip** After the desired process step is performed, the resist profile is to be removed. Resist removal can be carried out in a wet process using organic or inorganic solutions, for example, using thermal-assisted acid-based procedures, or by plasma stripping, which is the predominant technique in state-of-the-art projection lithography processes.

2.1 Contact and Proximity Lithography

Contact lithography, also called contact printing, is a lithographic pattern transfer technique that was used in the first years of IC fabrication (Lin, 1975; Heim, 1977; Henke et al., 1990).¹ A photomask is brought into contact with the resist. Light, typically coming from a broadband UV source, exposes the unmasked areas of the resist. The minimum achievable feature size is in principal determined by the feature sizes on the mask. However, in addition to limitations attributed to the photochemistry in the resist, the near-field diffraction effects lead to a rapid contrast loss when the distance of the image plane from the photomask increases, for example, at the bottom of thick resists. The most critical problem of contact printing is the formation of defects on the mask and the wafer caused by the surface contact.

As a consequence, contact printing was soon replaced by proximity printing in which a small gap of 10–50 nm between the photomask and the resist is maintained. Although avoiding mechanical contact and therefore significantly increasing productivity compared to the defect-prone contact printing approach, the near-field diffraction effects impose a harsh resolution limit on this approach. However, early ICs incorporated features larger than 100 nm. Moreover, with state-of-the-art mask aligners, features as small as 2–3 nm can be achieved, thus portraying a cost-efficient alternative for applications with limited resolution requirements such as micro-electromechanical systems (MEMS) or optical devices (Meliorisz et al., 2008).

2.2 Projection Lithography

Projection lithography employs a lens system that images the template, the photomask, onto the wafer. It was introduced in the 1970s and quickly displaced the then common contact and proximity printing approaches because it resolved many of the issues associated with them. In particular, the achievable throughput could be drastically increased. Moreover, today's projection systems demagnify the projected mask by a factor—typically four; leading to the common term *4x systems*. This property greatly eases the mask manufacture and inspection procedure, and gives rise to a higher yield.

The components of a lithographic projection system are schematically depicted in Figure 2.2: Light from a finite source is imaged by a condenser lens system to provide an uniform illumination of the photomask, which in its simplest form is a binary pattern consisting of transparent and opaque areas. A projection system generates an image of the mask in the photoresist on the wafer.

Projection lithography makes use of quasi-monochromatic light sources, that is, illumination sources that emit light whose wavelength (λ) varies only within a narrow bandwidth. The wavelength is one of the most significant variables for the resolution of the projection system. Table 2.1 gives a historic overview of illumination sources used in lithography and their wavelength. In early lithographic tools, a mercury arc lamp was employed as illumination source. At 436 nm, which is at the lower end of the visible spectrum, the mercury arc lamp shows a pronounced spectral power peak, termed *g-line*. This spectral line can be filtered and employed as a quasi-monochromatic source. Similarly, later tools applied another spectral line in the mid-ultraviolet (UV) regime, the *i-line*. In state-of-the-art lithographic systems, excimer² lasers are used. These types of lasers emit quasi-monochromatic light in the DUV spectrum by a stimulated generation of excited complexes of an inert gas and a halogen. Upon disassociation

¹Also see (Levinson, 2005, chap. 1).

²The term excimer is short for excited dimer. Since most of these types of lasers make rather use of excited complexes (exciplex), the term exciplex laser is more precise.




Figure 2.2: Components of the lithographic projection system: λ denotes the wavelength and σ the extent of the illumination source. The aperture angle (θ) of the projection lens and the refraction index (n) of the medium between the projection lens and the image plane define the NA.

Table 2.1: History of illumination sources in projection lithography (summarized from Mack, 2005). g-line and i-line denote the characteristic spectral lines at the lower visible and mid-ultraviolet end of the mercury arc lamp spectrum. The laser-induced plasma source listed at the bottom of the table will be used in EUVL, whose introduction is still uncertain.

type	wavelength	introduction
g-line of mercury arc lamp	436 nm	1978
i-line of mercury arc lamp	365 nm	1984
KrF excimer laser	248 nm	1987
ArF excimer laser	193 nm	1999
laser-plasma (EUVL)	13.5 nm	???

the excited complex (exciplex) emits laser light at a wavelength of 248 nm [Krypton fluoride (KrF)] or 193 nm [Argon fluoride (ArF)]. Wavelengths below 193 nm lead to less feasible optical conditions. For example, the long-anticipated transition to 157 nm, using a molecular fluorine laser (F_2), has been canceled due to a large number of technical issues mostly associated with optical and photochemical materials. Also the introduction of extreme-ultraviolet lithography (EUVL), which requires an all-catoptrics system, has been long-delayed and is still uncertain.

In projection lithography, the wafer is exposed progressively rather than in a single step, as opposed to contact and proximity printing. The two procedures shown in Figure 2.3 have prevailed: (1) In a *step-and-repeat* process, the entire chip is imaged onto the wafer in one step. Thus for each exposure, the movable wafer stage has to be positioned accordingly. The shutter is then opened, exposing one chip at a time (Figure 2.3(a)). (2) *Step-and-scan*, currently the predominant option, takes a lightly different approach. Here, the die is progressively exposed through a slit (Figure 2.3(b)). Although this technique requires both the wafer and the mask stage to be movable, which significantly complicates the design and the manufacture of lithographic tools, the improvements of the image quality, especially in terms of critical dimension (CD) uniformity, outweigh this disadvantage. The reason for the enhanced image performance can be mainly attributed to an aberration averaging due to the scanning procedure, but also to the fact that only a fraction of the lens area is utilized, allowing to actually minimize aberrations.

The lithographic imaging system depends on an ideal projection from the photomask onto the image plane. Unfortunately, because of inherent imaging imperfections, this cannot be achieved with a single lens. For example, spherical lenses show a different focus behavior for two parallel points depending on their distance to the optical axis of the lens. This is often referred to as spherical aberration. In addition to such inherent imperfections, there are various sources of wavefront aberrations including design and manufacturing issues, imperfect materials, or a deficient assembly or mounting of the lenses, including: *piston*, a deviation of the phase mean value, *tilt* of the wavefront, *defocus*, *coma*, an orientation-dependent distortion of off-axis points, or *astigmatism*, an orientation-dependence of the focal position. Because of the wavelength-dependence of the index of refraction, refractive optics are also affected by *chromatic* aberrations. The impact of aberrations on the lithographic process performance has been widely studied (see, for example, Brunner, 1997; Flagello et al., 1997; Otaka et al., 1997; Smith, 2000; Lalovic, 2001). In order to still facilitate imaging with low aberrations, a complex system of lenses is employed for the projection optics. Figure 2.4 presents an example of a fairly simple lens system for a 248-nm projection system with an NA of 0.8 (Ulrich et al., 2004, where also a historical overview on lenses is given; for a similar overview of another manufacturer see Matsuyama, 2006).

With these countermeasures, the lithographic projection system is indeed considered to be diffraction-limited. That is to say, diffraction is the most significant determining effect of the system. Although this assumption neglects all other actuating variables, it perfectly illustrates the main source for the resolution limit.




Figure 2.3: Exposure procedure of state-of-the-art projection tools: (a) Step-and-repeat: the entire die is exposed in one step; (b) Step-and-scan: the die is gradually exposed by a scanning slit.




Figure 2.4: Schematic set-up of a 248-nm projection lens system with an NA of 0.8 (from Ulrich et al., 2004). This fairly simple example demonstrates a design that combines a photorepeater and a microscope lens system into a projector for a step-and-scan system.

2.2.1 Resolution

One of the important properties of the optical system is the range of emitting angles of the projection lens. In lithography, this is defined in the same way as in microscopy, using the numerical aperture (NA):

$$\text{NA} := n \sin \theta, \quad (2.1)$$

where θ is the maximum collecting (half) angle of the lens (illustrated in Figure 2.2) and n is the refractive index of the medium below the lens. Traditionally, the lens is surrounded by air, which has a refractive index close to 1. The NA can be regarded as the “*collection efficiency*” of the projector. As such it is an important factor for the resolution limit of the projection system, which can be estimated through Rayleigh’s criterion (Rayleigh, 2004), originally introduced to rate the resolving power of telescopes: Given two stars separated by a small angle (with respect to the observation point). These stars can be regarded as two mutually incoherent point light sources. Thus the total intensity can be obtained as a sum of the intensities due to the individual source points.




Figure 2.5: Airy disk diffraction pattern set-up. Diffraction due to a circular aperture with radius a ; the distance between aperture and observation screen is f . An observation point can be defined by an angle ϕ and a radius δ on the screen or by the distance between the point and the aperture center (R), and the angle θ between the optical axis and the line between aperture center and observation point.

The point sources can be regarded as circular apertures whose Fourier transform—which corresponds to the Fraunhofer (far-field) diffraction formula³ is given by the Bessel function of the first kind and order. The resulting intensity distribution function is then

$$I(\theta) = I_0 \left(2 \frac{J_1(ka \sin \theta)}{ka \sin \theta} \right)^2, \quad (2.2)$$

³The diffraction far-field distance is given by the following condition:

$$F := \frac{a^2}{z\lambda} \ll 1,$$

where a is the size of the aperture (radius of the point source in this case), z is the distance from the aperture (point source) to the image plane and λ is the wavelength. F is called the Fresnel number.

where a is the radius of the aperture, θ is the angle between the optical axis and the line through the observation point and the center of the circular aperture, depicted in Figure 2.5, I_0 is the maximum intensity, and k is the wavenumber, with

$$k := 2\pi/\lambda.$$

The resulting distribution is called Airy pattern or Airy disk. It features a distinct intensity peak at its center and, moving outward, progressively declining fringes (see Figure 2.6).




Figure 2.6: Airy pattern.

The first root of the Bessel function—and hence of the Airy disk—occurs at an argument value of about 1.22π . According to Rayleigh's definition, the two stars are just resolvable if the maximum of the Airy pattern of one of them falls onto the first root of the other diffracted image. This situation is illustrated in Figure 2.7.




Figure 2.7: Airy pattern for two mutually incoherent point sources with an angular distance that, according to Rayleigh's criterion, renders them just resolvable in the far-field diffraction regime.

Therefrom, the resolution limit of two neighboring circular apertures is given by

$$\sin \theta = 1.22 \frac{\pi}{ka} = 1.22 \frac{\lambda}{2a} = 1.22 \frac{\lambda}{d}, \quad (2.3)$$

where d is the diameter of the aperture.

$$\sin \theta = \frac{\delta}{R},$$

with R being the observation distance and δ the distance on the screen resulting from θ (cf. Figure 2.5). For small angles θ , for which this analysis is performed, (2.3) a paraxial approximation can be applied:

$$\frac{\delta}{f} \approx \frac{\delta}{R} = 1.22 \frac{\lambda}{d}.$$

Because

$$\frac{f}{d} \approx \frac{1}{2\text{NA}},$$

for refractive indices close to one,

$$\delta = 0.61 \frac{\lambda}{\text{NA}}. \quad (2.4)$$

For several reasons the Rayleigh resolution does not fully adequately describe the resolution limit of the lithographic process, including:

- (1) For his criterion Rayleigh assumes point sources or circular apertures. The lithographic mask, however, consists of a multitude of different patterns with different sizes and characteristics.
 - (2) The properties of the photoresist and the related processing steps, none of which are captured in the Rayleigh assumption, significantly impact the behavior and thus the resolving capabilities of the overall lithographic process.
 - (3) Aberrations and other non-diffraction related variations are not considered.
 - (4) In contrast to immutable objects, as is the case with astronomic observations, lithographic illuminators and photomasks and can be altered to meet predefined characteristics (see Section 2.6).
 - (5) Subsequent IC process steps such as etching, deposition or implantation, are not taken into account.
- Despite these additional variables, the Rayleigh criterion serves as a viable qualitative estimation of the resolution limit of dense features and is frequently applied in lithography. Instead of a fixed factor that depends on the shape of the apertures, a variable factor, which incorporates technological aspects and which is hence called technology factor (k_1), is introduced:

$$r_{\min} = \delta = k_1 \frac{\lambda}{\text{NA}}, \quad (2.5)$$

where r_{\min} denotes the minimum half-pitch between two features. This measure can be used in both directions: to estimate the achievable size of dense features for a given k_1 factor or to assess how close a specific feature size is in proximity of the theoretical limit.

The ultimate theoretical resolution limit for a dense pattern is given in the case of two-beam interference, that is, when two plane waves arise from opposite propagation directions (achievable, for example, using off-axis illumination, which will be explained in Section 2.6.2). There, the minimum distance between features occurs for a path length difference between the two beams of half the wavelength, which in turn can be captured by the projection lens (of a given NA) for a k_1 value of 0.25.

2.2.2 Depth of Focus

Another important criterion for the projection system is its depth of focus (DOF). A displacement of the image plane from the focal length of the projection system leads to blur and thus to contrast loss in the image. DOF provides a measure for the maximum displacement such that the image exhibits a minimum resolution. A large depth of focus is essential as it stabilizes the lithography process with regard to variations caused, for example, by mechanical tolerances of both the wafer and the mask stage. More importantly, however, the photoresist is exposed over its entire thickness, for which an approximately constant image quality has to be maintained. Similarly to the Rayleigh resolution, the following equation gives an approximation for the depth of focus:

$$\Delta f := \pm \frac{\lambda}{2\text{NA}^2},$$

where the refractive index (n) is supposed to be 1.

Again, a technology factor, which takes the resist properties and resolution enhancement techniques into account, is introduced, giving:

$$\Delta f := k_2 \frac{\lambda}{\text{NA}^2}. \quad (2.6)$$

Because (2.6) uses a paraxial approximation, it becomes increasingly inaccurate with high NAs, leading to an overestimation of DOF (Lin, 2002). In addition, it is not valid for immersion lithography, which we shall introduce in Section 2.2.8. There, the DOF is underestimated, as the refractive index (of the immersion fluid) is accounted for twice, since

$$\lambda = \lambda_0/n \quad \text{and} \quad \text{NA} = n \sin \theta_0.$$

Lin (2002) thus derived the following modified DOF formula, which he showed provides a more accurate estimate for high-NA and immersion lithography:

$$\Delta f := k_3 \frac{\lambda_n}{\sin^2(\theta_n/2)}, \quad (2.7)$$

where λ_n is the wavelength in the imaging medium, and k_3 is a technology factor. The angle θ_n is the emerging angle from the lens corresponding to one pitch on the mask.

2.2.3 Lithographic Process Evaluation

The previous DOF considerations are suited for a general assessment of a given lithography option. In order to evaluate an actual process, a more practice-oriented methodology is commonly applied, which can be conducted experimentally or numerically. It relates variations of the focal plane to the CD of a given feature. In general, the CD is the size of a transferred pattern, for example, printed in the photoresist or obtained after subsequent process steps such as etch or deposition. The CD can also be measured in the aerial image (Figure 2.8(a)). There a threshold intensity at which the pattern prints is assumed. The CD in resist profiles can be evaluated at different height positions (Figure 2.8(b)). Most commonly, the bottom CD, which is measured at ten percent of the height of the resist profile, is used.

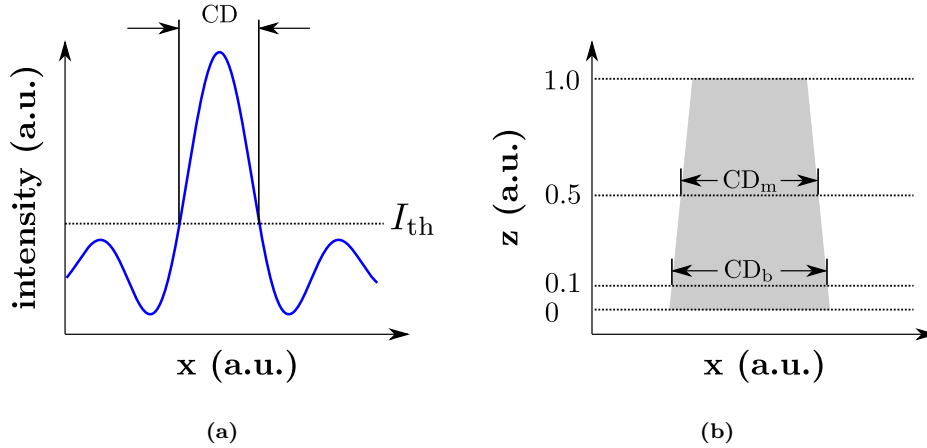


Figure 2.8: Critical dimension measured (a) using a threshold intensity (I_{th}) in the aerial image, or (b) at the bottom (10 percent of the height; CD_b) or at the middle (50 percent of the height; CD_m) of the photoresist profile.

Now the focal range in which the CD is still on target can be determined, e.g., by a brute-force simulation sequence. In order to account for acceptable size variations, a ± 10 percent CD tolerance

is often allowed for. The obtained focal range then defines the depth of focus for one or more given features.

In a similar fashion, we can determine another factor that strongly impacts the stability of the lithographic process, the exposure latitude. For a plain aerial image analysis, the intensity threshold is varied such that the CD is on target, again allowing for a predefined (e.g. ± 10 percent) CD tolerance. In case the photoresist is taken into account, the CD is measured at the according position in the resist profile. Typically, the exposure latitude is given as a relative value, i.e.:

$$\Delta E := \left| \frac{E(0.9 \text{ CD}) - E(1.1 \text{ CD})}{E(\text{CD})} \right|$$

There are a number of different schemes to visually demonstrate the relation between CD, exposure latitude and depth of focus (Bossung, 1977). To determine the lithographic process window (PW), that is, the domain of a stable process, the following technique is well established (illustrated in Figure 2.9(a)): The maximum and the minimum admissible CDs are plotted as a function of both the focus and the dose—or, in the case of an aerial image analysis, the threshold intensity.⁴ A rectangle or an ellipse fitted into the area between both curves constitutes the feasibility region of the process for the given feature. A rectangle is used to account for systematic errors while an ellipse is suited to depict random errors. A similar, typical analysis is shown in Figure 2.9(b). There the dose latitude is depicted as a function of the maximum depth of focus, using either a rectangle or circle fit.




Figure 2.9: Lithographic process window: (a) Rectangle and circle fit into region encompassing ± 10 percent of the target CD. (b) Dose latitude as function of depth of focus using either fit as demonstrated in (a).

It is often reasonable for process windows to be based on further parameters in addition to the CD. Common objectives are side-wall angles of the resist profile and the total resist height.

2.2.4 Illumination Properties

The previous analyses assume that the illumination source is totally spatially coherent. Given an incoherent light source, that is the case if the light incident on the mask originates from one point source. In

⁴As Fuard et al. (2002) argue, the correspondence between dose and intensity threshold is more adequately stated by:

$$\text{th} := \frac{a}{d - b},$$

where th is the threshold, d is the dose, and a and b are process-related empirical constants. For simplicity, this methodology is not used in this example, instead threshold and dose are simply identified.

optical lithography, however, extended sources illuminating the mask from a range of different incidence angles are applied. The main advantages of this type of illumination source is its ability to image objects with higher spatial frequencies and a reduction of interference-induced sidelobes. To conveniently illustrate this property, the notion of diffraction orders is briefly introduced in the following paragraph.⁵

Given a periodic grid of slits, each of which has an infinite height and a small width compared with the wavelength. Let these slits be omitted (i.e., transparent) from an otherwise opaque plane. This pattern is coherently illuminated by a plane wave imaging it onto a screen in the diffraction far-field. Because of the infinitesimal width of the slits we can treat this problem as two-dimensional. The intensity function for a two-dimensional slit in the Fraunhofer approximation is given by (the Fourier transform of the rectangular function):

$$I(\theta) = I_0 \operatorname{sinc}^2 \frac{ks \sin \theta}{2}, \quad (2.8)$$

where s is the slit width, k is the wave number, θ is the angle between the optical axis and the line from the slit to the observation point, and I_0 is the intensity in the center of the imaged pattern. The interference function for a grating with N slits is as follows:

$$I(\theta) = I_0 \frac{\sin^2 \left(\frac{Nkd \sin \theta}{2} \right)}{\sin^2 \left(\frac{kd \sin \theta}{2} \right)}, \quad (2.9)$$

where d is the distance between two slits. The flux density distribution is thus obtained by a multiplying the intensity and the interference functions:

$$I(\theta) = I_0 \operatorname{sinc}^2 \frac{ks \sin \theta}{2} \frac{\sin^2 \left(\frac{Nkd \sin \theta}{2} \right)}{\sin^2 \left(\frac{kd \sin \theta}{2} \right)}. \quad (2.10)$$

The graphs of the intensity, the interference and the flux density distribution function for a grating with seven slits, where the distance between the slits is $\lambda/2$ and the slit width is $\lambda/16$ is shown in Figure 2.10.

The interference function (2.9)—and hence the flux density distribution (2.10)—exhibits maxima when

$$\sin^2 \left(\frac{kd \sin \theta}{2} \right) = 0.$$

Thus

$$\frac{kd \sin \theta}{2} \in \{m\pi \mid m \in \mathbb{Z}\},$$

yielding maxima of (2.10) when

$$\sin \theta \in \{m\lambda/d \mid m \in \mathbb{Z}\}.$$

The integer m is referred to as diffraction orders, since it can be considered the factor of the beam decomposition by the grating. Figure 2.11a illustrates the concept of diffraction orders.

By extending the source and thus by introducing additional incoherent point sources, additional light directions are propagated through the grating (see Figure 2.11b), which in the resulting superposed image leads to an increased contrast.

The extend to which the illumination source is either coherent or incoherent is termed *partial coherence*. It can be expressed as the ratio between the aperture angle of the source and the projection lens (see Figure 2.12):

$$\frac{\sin \beta}{\sin \alpha}$$

Thus a partial coherence factor of zero renders the illumination source fully coherent, whereas a factor of infinity would yield a totally incoherent source. In optical lithography imaging systems, extended sources are realized using Köhler illumination (Köhler, 1893). This method has been developed by

⁵The derivation broadly follows Born and Wolf (1999, Section 8.6.1, pp. 446–453) and Hecht (2002, Section 10.2.3, pp. 460–464).




Figure 2.10: Flux density distribution of a two-dimensional grating with seven slits, with a slit distance (d) of $\lambda/2$ and a slit width (s) of $\lambda/16$. The intensities of the intensity function [(2.8); thin solid line], the interference function [(2.9); dotted line] and the flux density distribution function (2.10) are normalized. Intensity peaks emerge at multiples of λ/d , called diffraction orders.

August Köhler in 1893, in order to eliminate deficiencies of former illumination sources in microscopy. The basic idea is to introduce a condenser lens which produces an image of the illumination source. Thus brightness irregularities, which in microscopy often led to filament images, are compensated for.

2.2.5 Mask Error Enhancement Factor

One of the main technical advantages of projection lithography is its demagnification potential, which allows for the use of larger photomasks, both facilitating the manufacture of reticles and hence reducing a key source of errors. Indeed, defects on a mask for an $M\times$ system (a system with a demagnification of factor M) lead to an M times smaller imaging error. While a rather linear relation between mask line widths and printed CDs occurs in the scope of larger features, this property is not valid for smaller features, especially when close to the resolution limit. Figure 2.13 depicts this behavior. The same figure also shows the significant difference between the characteristics of a dense, periodic structure and an isolated line. The following simulation settings are applied: 193 nm wavelength, circular illumination with a partial coherence factor (σ) of 0.6 and an NA of 0.8. An intensity threshold approximately optimal for medium-sized dense lines/spaces is assumed.

A metric for the this non-linear impact of mask errors on the printed CD was first proposed by Maurer (1996) and often referred to as the mask error enhancement factor (MEEF). Later refined by Iba et al. (1995) by taking the slope of the proximity curve into account, the there called mask error factor (MEF) is defined at a fixed pitch as:

$$\text{MEF} := \frac{\partial \text{CD}_w}{\partial \text{CD}_m}, \quad (2.11)$$

where CD_m is the feature size on the mask (in wafer scale) and CD_w is the corresponding size on the wafer. According to this metric, features with a mask error factor close to one can be considered well-natured, whereas for large factors, a sound CD control is more intricate to achieve.

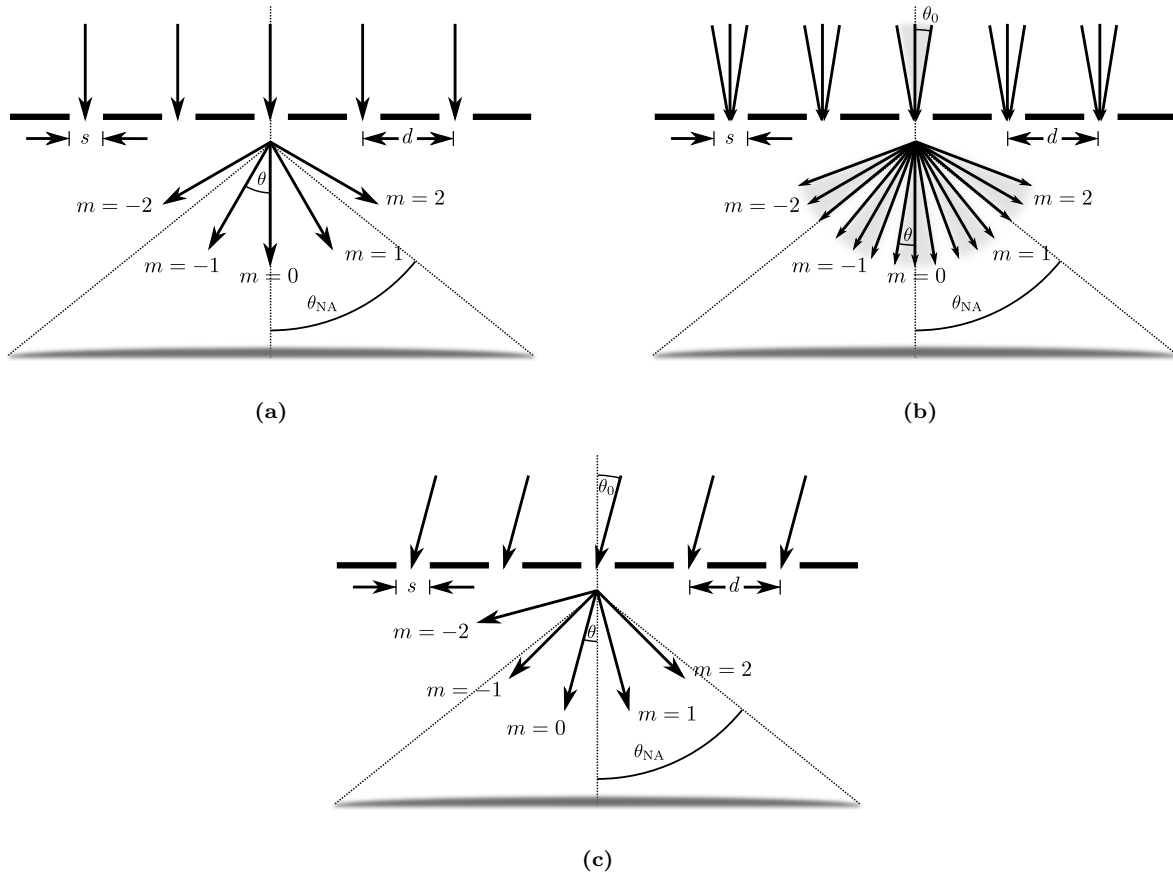


Figure 2.11: A two-dimensional regular grating, with a distance d between the slits and a slit width of s . (a) A plane wave incident on the grating results the three diffraction orders (m , separation angle θ) $-1, 0$ and 1 that are collectible with the given NA—illustrated by the collection half-angle θ_{NA} . (b) By introducing an extended source and thus a range of incidence angles (θ_0), the number of diffraction orders within the collecting angle of the lens is increased, thus extenuating the minimum resolution. (c) Using oblique illumination, the diffraction orders are shifted. In the depicted extreme case, an additional order is so shifted into the NA range of the lens.




Figure 2.12: The partial coherence factor (filling factor) is given by the size ratio between source extension and projection lens.

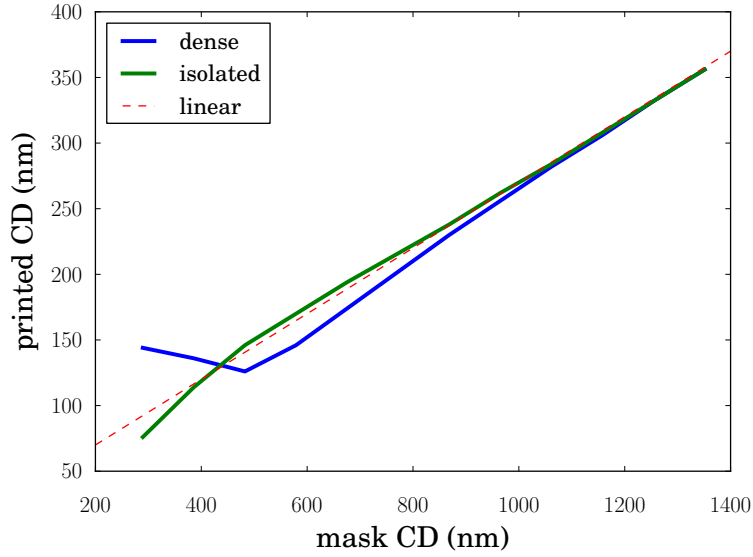


Figure 2.13: Printed line width as a function of the line width on the mask. For larger widths, target and mask CDs show a linear relation. This linearity is not maintained for smaller CDs. Moreover, dense and isolated features show a significantly different behavior for CDs close to the resolution limit.

Rigorous investigations on the causes of large MEEFs and metrics relating MEEF, defocus and exposure doses were presented by Wong et al. (2000). A matrix notation taking different mask perturbation scenarios into account was developed by Granik and Cobb (2002) and applied to optical proximity correction (also see Cobb and Granik, 2002).

2.2.6 Contrast and Modulation Transfer Function

When imaging dense features such as a lines/spaces pattern, the performance of the projection system can be adequately measured by evaluating its contrast, which is defined as:

$$C := \frac{I_{\max} - I_{\min}}{I_{\max} + I_{\min}}, \quad (2.12)$$

where I_{\max} and I_{\min} are the maximum and the minimum intensities of the output image, respectively.

A more rigorous measure is constituted by the modulation transfer function (MTF), which correlates the spatial frequency of features and the obtained contrast, and additionally takes the partial coherence factor into account. The MTF can be directly derived from the transmission cross-coefficients of the Hopkins imaging formulation, which will be introduced in Section 3.2.6. For the extreme cases—full coherence or incoherence—the MTF is defined as follows:

$$\mathcal{M}_{\text{coherent}} := \begin{cases} 1 & \nu \leq \text{NA}/\lambda \\ 0 & \text{otherwise} \end{cases}$$

$$\mathcal{M}_{\text{incoherent}} := \begin{cases} \frac{2}{\pi} \left(\arccos \frac{\lambda \nu}{2\text{NA}} - \frac{\lambda \nu}{2\text{NA}} \sqrt{1 - \left(\frac{\lambda \nu}{2\text{NA}} \right)^2} \right) & \nu \leq 2\text{NA}/\lambda \\ 0 & \text{otherwise} \end{cases},$$

where ν is the spatial frequency of features.

As shown in Figure 2.14, the coherent case ($\sigma : 0$) provides a maximum contrast for isolated features (having a low spatial frequency), but becomes zero at frequencies larger than λ/NA . In the incoherent

case ($\sigma : 1$), on the other hand, a contrast is maintained up to normalized frequency values of two. For isolated features, however, the contrast is significantly lower in comparison with the coherent illumination. Compromise scenarios can be obtained by setting the partial coherence factors to values in between 0 and 1.

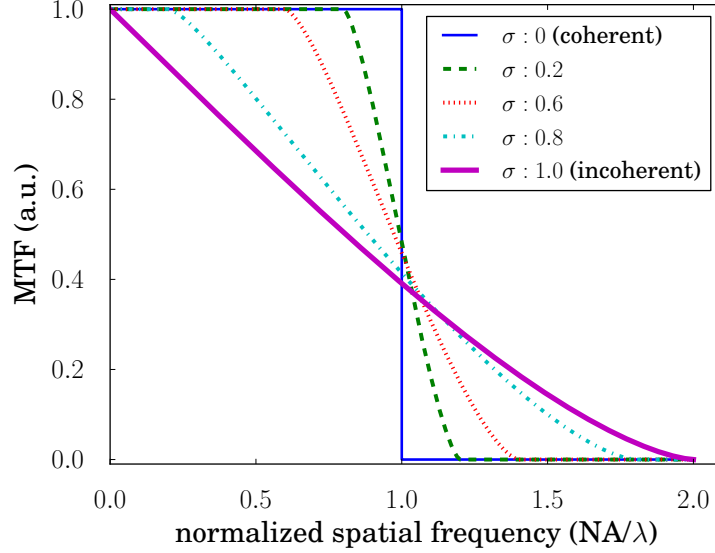


Figure 2.14: Modulation transfer function for different partial coherence factors.

2.2.7 Normalized Image Log Slope

The image log slope (ILS) captures the gradient of the image at edges, that is, the transitions of a dark to a bright region—or vice versa. For a one-dimensional image, the ILS at a point x is defined as the normalized intensity gradient:

$$L := \frac{1}{I} \frac{dI}{dx} = \frac{d \ln I}{dx}. \quad (2.13)$$

In order to relate this metric to the exposure latitude, it is often useful to scale the slope metric with the line width of the measured feature w :

$$N := w \frac{1}{I} \frac{dI}{dx} = w \frac{d \ln I}{dx}. \quad (2.14)$$

The so obtained measure, termed normalized image log slope (NILS), can be considered a local contrast metric, which is often applied to monitor the impact of dose variations on critical features.

2.2.8 Immersion Lithography

In immersion lithography, the medium between the projection lens and the wafer is changed from air to a fluid, thus increasing the achievable effective numerical aperture by the index of refraction of this immersion fluid (2.1). This allows for the exploitation of higher diffraction orders so that in turn both the resolution and the depth-of-focus of the projection system are greatly improved (Lin, 2004). The principle is demonstrated in Figure 2.15: In contrast to dry lithography, where the gap between the last lens and the top layer of the wafer (e.g., the resist or an anti-reflective coating) is filled with air, a fluid is constantly dispensed between lens and wafer, and recovered on the opposite of the wafer motion direction. This is done mostly to reduce the contamination load.




Figure 2.15: Schematic comparison of wafer stage for (a) dry and (b) immersion lithography. In the immersion case, a blob of purified and degassed water (or another immersion fluid) is applied between the last lens of the projection system and the resist.

Although immersion has been proposed and used in microscopy for more than 100 years (Abbe, 1873), it has only been recently introduced in lithography. The most critical technical issue is the wafer scanning motion in spite of which the immersion fluid has to be maintained between the last lens and the wafer. In addition, the fluid has to meet a number of opto-chemical requirements: While the refraction index of the immersion fluid should be maximized, in contradiction, a low absorbance is to be achieved at the same time. Furthermore, the fluid has to be compatible with both the last lens and the resist material.

State-of-the-art projection systems employ purified degassed water as the immersion fluid (Smith et al., 2004), enabling the 32 nm HP node.⁶ At 193 nm, water has a refraction index of about 1.44. The practically achievable effective NA is about 1.35.⁷ A major concern with 193 nm water immersion (193i), is the immersion-specific defects (Okoroanyanwu et al., 2005). Especially water bubbles caused by the water dispensing at high scanning speeds (de Bisschop et al., 2004; Fan et al., 2004; Kawai et al., 2006; Kocsis et al., 2004; Switkes et al., 2005) and the interaction of the immersion fluid and the photoresist (Taylor et al., 2004) posed practical problems, all of which can be considered to be solved.

Although suited for the 32 nm process, even higher NAs could be conducive for succeeding nodes (Zhou et al., 2007). For that purpose, high-index immersion fluids (French et al., 2007; Peng et al., 2005), lens materials (Parthier et al., 2007; Burnett et al., 2005) and photoresists (Matsumoto et al., 2008) are required. In addition to material-related issues (Zimmerman et al., 2007; Sewell et al., 2007), the hyper-NA regime results polarization-dependent artifacts (Flagello et al., 2005; McIntyre et al., 2008) that need to be considered in the process design.

2.3 Extreme-Ultraviolet Lithography

Operating at a wavelength of 13.5 nm or below, EUVL uses all-reflective optics. Thus both the condenser and the projection lenses have to be realized as mirrors. In addition the photomask blank has to be highly reflective. Figure 2.16 demonstrates the principal set-up of a state-of-the-art EUV projection system. Commonly a laser-produced plasma is used. Alternatives including a discharge-produced plasma source have also been proposed. In contrast to the laser sources used in optical lithography, light emitted by

⁶Some manufacturers have employed immersion lithography already for 45 nm HP.

⁷This corresponds to $\sin(\theta)$ of 0.95.

the source has to be focused before it can be propagated through all-reflective condenser system. A further mirror collects the light from the condenser and reflects it onto the reticle, which is composed of a substrate, a multilayer mirror and an absorber. The light reflected from the reticle is collected and transferred through projection system. In the example shown in the figure, a six-mirror system is depicted, which is suited for NAs up to 0.35. In EUVL, all optical systems have to be placed in vacuum. In order to avoid mutual contaminations and to simplify the handling of the individual components, the stepper is typically divided into several vacuum chambers, for the illumination optics, the projection system, and the reticle and the wafer stage.




Figure 2.16: EUVL projection system employing all-reflective optical components. In this example, a six-mirror projection system suited for NAs of up to 0.35 is depicted. In order to reduce down-times in case of maintenance, condenser, projector, and the mask and wafer stages are placed into individual vacuum chambers.

Of course, also the plasma source, which is one of the critical components of an EUV system (Banine et al., 2000), is contained in a vacuum housing. It is moreover decisively isolated from the condenser optics to prohibit contaminations from the plasma. laser-produced plasma (LPP) source is one of the candidates for early production EUV steppers. Figure 2.17 illustrates the major components of tin-fueled LPP source (Koay et al., 2003; Endo et al., 2007; Brandt et al., 2009): The primary source is a high-power laser, today commonly a CO_2 laser. The laser beam is incident on a tin droplet—ejected from a droplet generator—producing the plasma. The emitted photons are collected and intermediately focused. The intermediate focus (IF) constitutes the interface between the source and the illumination optics. Solid plasma targets lead to a significant amount of debris. Especially the collector mirror—constituting a Faraday cup—is highly affected by this debris, in that it drastically reduces the uptime and the lifetime of the collector. One countermeasure is the application of a magnetic field that serves as plasma guide (Niimi et al., 2003; Harilal et al., 2005). The use of ambient gas has also been proposed (Harilal et al., 2006). Using gaseous plasma targets such as xenon (Shields et al., 2002) reduces debris but also decreases the conversion efficiency (Koay et al., 2003), that is, the ratio of the input energy of the plasma producing source—e.g., the laser—versus EUV irradiance within a two percent bandwidth tolerance around the target wavelength (e.g., 13.5 nm)⁸.

⁸For a rigid definition of the term *conversion efficiency* see for example Bakshi (2006)




Figure 2.17: Example of a laser-produced-plasma sourced used in EUVL: A high-power pulsed laser (e.g. CO₂ laser) on a tin droplet target, which is ionized and so emits EUV photons. A mirror collects the photons and directs them towards an intermediate focus (IF). An electro-magnetic field is employed in order to mitigate the impact of debris on the collector mirror.

As an alternative to LPP sources, plasma that emits EUV light can also be produced by a gas discharge. One approach is to generate a gas jet of the plasma fuel—such as xenon—which is pre-ionized and finally ignited by a high-current pulsed power supply (e.g., see Stamm et al., 2003; McGeoch and Pike, 2003; Kishi et al., 2008). A high-speed gas curtain (e.g., using helium) is applied for debris mitigation. Although otherwise ideal as a plasma target, the EUV light emitted from xenon shows a peak at a wavelength of about 11 nm. Tin on the other hand has an EUV emission peak that is very close to the designated wavelength of 13.5 nm. Kishi et al. (2008) propose to trigger the discharge by an neodymium-doped yttrium aluminium garnet (Nd:YAG) laser. As an alternative, the plasma can also be triggered by a hollow cathode (Pankert et al., 2003). Other approaches employ a dense plasma focus (DPF) device (Partlo et al., 2000) or a capillary discharge (Klosner and Silfvast, 1998). Like in the LPP case, metallic targets such as tin can drastically improve the conversion efficiency but also increase the debris load (Teramoto et al., 2007).

With a conversion efficiency of less than 4 percent, seen in pre-production steppers shipped today, both the condenser and the projection optics but also the reticle blanks have to exhibit an extremely high peak reflectivity in order to render a viable next-generation lithography (NGL) candidate. This is achieved by a Bragg reflector (Figure 2.18 shows such an interference multilayer mirror underneath the absorber). More specifically, a 40-layer MoSi multilayer—with a layer thickness of approximately 4.11 nm (Mo) and 2.74 nm (Si), respectively—has proved to provide an high peak reflectance (of up to 70 percent) at the 13.5-nm EUV wavelength. The performance of the multilayer mirrors strongly depends on the absence of defects. For example, defects that are buried in the multilayer or below on the substrate are of major concern, and their impact has hence been intensively studied (Pistor et al., 2000a; Ito et al., 2001; Gullikson et al., 2002; Evanschitzky et al., 2003; Stearns et al., 2004; Farys et al., 2005; Lam and Neureuther, 2005; Clifford and Neureuther, 2009; Huh et al., 2010). Several inspection (Tezuka et al., 2004; Barty et al., 2005; Hamamoto et al., 2005; Goldberg et al., 2006; Tanaka, 2006; Wurm et al., 2007; Lin et al., 2008) and repair (Yan et al., 1998; Mirkarimi et al., 2002; Barty et al., 2003; Liang et al., 2005; Liang and Gullikson, 2008; Clifford et al., 2010) approaches have been proposed.

In addition, the high absorbance of the multilayer mirrors at the EUV wavelength leads to a high heat load that in turn affects the stability of the MoSi mirrors. For example, at the interface between the molybdenum and silicon, an amorphous intermixing layer is formed. This effect is more severe under annealing condition similar to the heat load under EUV radiation (Stearns et al., 1990; Nedelcu, 2007).

It can be counteracted by interlayer barriers or by using molybdenum or silicon compounds instead of MoSi (Singh and Braat, 2000; Feigl et al., 2001; Braun et al., 2002; Bajt and Stearns, 2005; Yulin et al., 2006). Moreover, the thermal load also leads to topographic instabilities such as deformations (Chiba et al., 2002). Another issue is the oxidation of the surface layers of the mirror (Benoit et al., 2005). Protection layers, called capping layers, are commonly introduced to decrease the oxidation rate (Bajt et al., 2003). Especially the collector and the condenser optics are intensely subjected to heat and thus suffer higher damage than the components of the optical system.




Figure 2.18: EUV photomask with a substrate, a reflective multilayer consisting of molybdenum-silicon bilayers (typically 40–50), and an absorber. In order to decouple incident and reflected light, a small inclination angle (θ) of typically 6° is applied. The dashed line in the sketch indicates the normal direction, the incidence due to the angle θ is shown by the solid line.

Figure 2.18 shows a schematic of an EUV mask. The MoSi multilayer mirror is placed on top of a low thermal expansion material (LTEM) substrate. Typically a silicon or ruthenium capping layer, mainly to protect the molybdenum from oxidation, is deposited between the multilayer system and the absorber. The absorber commonly consists of two layers: (1) The actual (conductive) absorber—for which a wide range of materials has been proposed (Yan et al., 1999; Mangat et al., 2000; Robic et al., 2002). (2) And, beneath, a buffer layer that has both the role of an etch stop layer and a protection layer for absorber repair processes (Yan et al., 1999; Lee et al., 2002). In order to decouple the incident and the reflected light, a small inclination angle (θ) of typically 6° is applied. Thus telecentricity errors are inherent in EUV lithography and lead to a significant orientation dependence of the pattern placement and the line widths. Since these errors are focus dependent, they induce focus shifts and thus lead to process window asymmetries (Bollepalli et al., 1998; Otaki, 2000; Krautschik et al., 2001; Chen, 2003; Besacier and Schiavone, 2004). Yan (2002) proposed a phase-error interpretations of these effects. Bollepalli further demonstrated a significant imaging bias between features on dark and bright field masks. Cerrina et al. (2000) showed that the telecentricity errors, also called shadowing effect, increase with higher NAs. Systematic simulation surveys of the impact of the EUV mask absorber properties have been performed by Schiavone et al. (2001); Robic et al. (2002); Erdmann et al. (2003a). Nakajima et al. (2008) studied EUV mask artifacts under the influence of scanner aberrations, employing a Zernike sensitivity analysis. Interpreting the mask artifacts as aberrations, the authors introduced the concept of an aberration budget for EUV. Elsewhere (Erdmann et al., 2010) we have presented a rigorous analysis of the wavefront diffracted from an EUV mask. Similarly to the former authors, we have incorporated an aberration analysis based on the Zernike formalism in order to quantify the impact of mask-induced artifacts on the imaging performance. When extended to higher order Zernike terms, such a technique could be applied to optical proximity correction (OPC) models or even become an integral component of future source/mask optimization approaches.

Another problem of EUVL is the availability of appropriate resist materials. On the one hand, the resist is to be highly sensitive in order to allow for a high wafer throughput—especially in view of the limited conversion efficiency of the source (Kozawa et al., 2006). On the other hand, due to shot noise, too

sensitive a resist is substantially more prone to line-edge roughness. This effect is additionally aggravated by low-energy secondary electrons that can lead to unintentional acid deprotections (Kozawa et al., 2007). The high absorption of the resist material in the EUV regime may also result pattern instabilities (Jouve et al., 2005; Chang et al., 2007). Moreover, the vacuum environment boosts outgassing of resist components, leading to contaminations (Hada et al., 2004; Kobayashi et al., 2008).

Because of the failure to be timely available, EUV has to aim at considerably smaller half pitch nodes than originally planned. Thus different strategies—including advanced resolution enhancement techniques but also even smaller wavelength (Otsuka et al., 2010; Sasaki et al., 2010; Wagner and Harned, 2010)—are currently evaluated in order for EUV to become the optical lithography successor.

2.4 Electron-Beam Lithography

Electron beam lithography (often abbreviate e-beam lithography) is a direct writing technique (Levinson, 2005, chap. 7, pp. 252; chap. 12, pp. 398). That is, in contrast to most other lithography approaches it involves no template. Instead an electron beam is scanned over a resist coated surface. This operation is comparable with that of a scanning electron microscope (SEM), from which this technique is actually derived. Because of its high flexibility—due to the fact that it is maskless—and its superior resolution in the nanometer regime, e-beam lithography plays a very important role for prototyping applications and in the manufacturing of niche products. The significant lower throughput compared to other lithography approaches, which is caused by its scanning nature, renders e-beam lithography in general not suitable for high-volume manufacture (HVM) scenarios. Electron beam lithography still plays a very important role for IC fabrication, for example, to study the physical and electrical behavior of future IC generations. In addition, it is by far the predominant patterning option for photomasks in lithography techniques that involve masks.

Like photolithography, e-beam lithography makes use of photoresists, including positive and negative, and chemically amplified and non-amplified materials. In contrast to photonic lithography, e-beam lithography is not diffraction-limited. Proximity effects, which are also present under e-beam exposure, are conveyed by electron scattering, of which there are different mechanisms such as backscattering—either from the substrate or through electron-nucleus interaction in the resist—and forward scattering through electron-electron interaction (Chang, 1975; Parikh, 1979). In addition, much like in EUV lithography, a major part of the resist exposure is by dint of secondary electrons, so contributing to proximity artifacts.

Further miniaturization will dramatically increase the cost of ownership associated with lithography. One of the critical factors is the fabrication cost of reticles. A major improvement of throughput would hence render maskless techniques such as e-beam lithography an attractive NGL candidate. One approach is the implementation of projection e-beam lithography by utilizing beam shaping techniques (e.g., see Pfeiffer, 1978). An alternative solution is the use of parallel e-beams (Lischke et al., 1989). Advances in both fields have been currently reported (Petric et al., 2009; Gramss et al., 2010; Klein et al., 2010; Lam et al., 2010; Wieland et al., 2010; Yamada et al., 2010). In addition to aiming at an EUV or double patterning⁹ alternative, parallel or shaped e-beam approaches are of course also qualified to speed up the write time of photomasks, which would help to reduce the cost of ownership of EUV and especially of double patterning equipment.

2.5 Alternative Lithography Approaches

Similarly to e-beam lithography, *focused ion beam (FIB) lithography* is a direct write technique (Seliger and Fleming, 1974; Melngailis, 1993). Because of lower proximity effects and highly ion-sensitive resist materials, FIB provides a much improved resolution compared to e-beam lithography. Until now

⁹Used here as an umbrella term for techniques that involve multiple exposure steps (and therefore also multiple photomasks). See Section 2.6.4 for a detailed discussion.

though, ion-beam lithography lacks a durable and reliable source with a sufficiently high yield. Ion-beam lithography can also be implemented as a resist-less approach. In that case, it is applied as a direct sputtering approach, requiring an even higher ion yield, and has thus a reduced throughput. Another variant, *ion beam projection lithography (IPL)*, makes use of a condenser optics. This technique allows for larger areas to be patterned in a single exposure. It can be implemented either using masks (Randall, 1983; Bartelt, 1986) or in a maskless approach (Stengl et al., 1987), which resembles the shaped-beam technique used in e-beam lithography. Multiple ion-beam techniques with an almost identical implementation as parallel e-beam approaches have been proposed recently (Platzgummer et al., 2007; Platzgummer, 2010).

Another mask-less approach is the *direct laser writing lithography*, in which a focused laser beam exposes a photosensitive material. This technique plays a role—albeit a secondary one compared to e-beam lithography—in the fabrication of photolithographic reticles (Mac Donald et al., 1984)¹⁰. Later, photomask repair techniques were also proposed (Lieberman et al., 1996; Haight et al., 1998). Like other direct writing techniques, direct laser writing suffers from slow throughput and is thus not well suited for high-volume chip manufacture; but it is a well-established technique for the fabrication of gratings, holograms and photonic crystals (e.g., see Mashev and Tonchev, 1981; Deubel et al., 2004, see) [also see references for *holographic lithography* below], with a wide range of applications including the generation of diffractive optical elements (Koronkevich et al., 1984; Björnängen et al., 2004), used, for example, in state-of-the-art lithographic illuminators (see Section 2.6.2). In addition, direct laser writing may present an economic alternative for tailored or low-volume products (e.g. Schomburg et al., 1997; Fleury et al., 1995). With the recently demonstrated employment of a continuous wave laser (Thiel et al., 2010), direct laser write may attain increasing popularity.

In *nanoimprint lithography* a template is brought into mechanical contact with a monomer or polymer (Michel, 1992). The pattern is hence virtually stamped into the resist, which is then hardened by heating or through UV radiation. Potentially, nanoimprint features a medium to high throughput and relatively low mask and tool costs. There are, however, several technological issues including defects, overlay issues and mask lifetimes that currently prevent nanoimprint from becoming an NGL candidate to be reckoned with. Still, nanoimprint may play a more important role in low-volume manufacture (LVM) such as the fabrication of lithographic masks or nano-electromechanical systems (NEMS) in the future.

Other alternatives include:

- *Atom lithography* (Berggren et al., 1995; Timp et al., 1992; Oberthaler and Pfau, 2003),
- *x-ray lithography* (Spears and Smith, 1972; Feder et al., 1977; Neureuther, 1980; Silverman, 1998)
- *interference lithography* (Hinsberg et al., 1998; Hoffnagle et al., 1999; Raub et al., 2004), *interference lithography* using (soft) x-rays (Solak et al., 1999, 2002), *interference-assisted lithography* (Fritze et al., 2005; Greenway et al., 2008), *holographic lithography* (Beesley and Castledine, 1970; Campbell et al., 2000; Moon et al., 2004; Xie et al., 2007),
- *scanning probe lithography* (McCord and Pease, 1986; Betzig and Trautman, 1992; Snow and Campbell, 1994; Salaita et al., 2006; Piner, 1999),
- *near-field lithography* (Pohl et al., 1984; Minne et al., 1998; Alkaisi et al., 1999; Ghislain et al., 1999; Goodberlet, 2000; Kik et al., 2001; Blaikie et al., 2006; Srituravanich et al., 2008), and
- *zone plate array lithography* (Gil et al., 2003).

2.6 Resolution Enhancement Techniques

The quest of the semi-conductor industry for increasingly small feature and pitch sizes prompts for a constant reduction of the resolution limit in lithographic projection systems. The most obvious means is the reduction of the wavelength or an increase of the numerical aperture of the projection lens (2.5). There are hard constraints to both of these methods: The introduction of new wafer steppers using a smaller wavelength involves an enormous technological effort. Not only the lithographic appliances have to be re-designed and produced, also materials such as lenses and photoresists have to be adapted to the new scenario. For example, the introduction of extreme ultra-violet light, which involves the conversion

¹⁰A comprehensive overview on direct laser writing was authored by Allen (2002).

to a reflective projection system, has been postponed several times due to technical issues such as the illumination source, mask blank defects and heating in the mirror (lens) system.

Just as well, the NA is subject to practical constraints including lens design issues. NAs close to the theoretical limit of 1.0 (in free space), which can hardly be realized, lead to extreme incidence angles that in turn compromise the lithographic imaging performance (Brunner et al., 2002). One way to overcome this limitation is the introduction of immersion lithography (see Section 2.2.8), whose NA increasing capability is again limited, by the index of refraction of the immersion fluid.

There are, however, a number of further actuating parameters that influence the lithographic imaging performance, most importantly related to photomask and illumination source settings. The involved techniques, which are called resolution enhancement techniques (RET), can be considered the foundation for further extensions of optical and EUV lithography beyond current limits.

2.6.1 Optical Proximity Correction

Due to the band-limiting property of the projection lens, the image produced by the lithographic system differs from the objective, the photomask, as high frequencies are lost. This effect increases with decreasing pitch sizes.

Image degradation caused by proximity can be classified into the following four different categories (Cobb, 1998; Wong, 2001):

- *Iso/dense bias*: The number of diffraction orders is determined by the pitch of elements on the mask. Dense arrangements lead to less orders than isolated ones. The number of contributing orders significantly impacts the image. Thus an identical feature will create a different pattern depending on its proximity characteristics. For that reason, Wong terms this effect proximity effect. Figure 2.19 portrays the CD of a line/space pattern as a function of the duty ratio, that is, the ratio between the line width and the width of the adjacent space. The target CD is 97 nm. A wavelength of 193 nm, an NA of 0.8 (dry) and an annular illumination (inner radius: 0.2, outer radius: 0.4) were used. The CD was evaluated at a threshold of 0.3. The leap at the duty ratio of 1:1.25 and 1:1.5 can be attributed to additional diffraction order that then contribute.

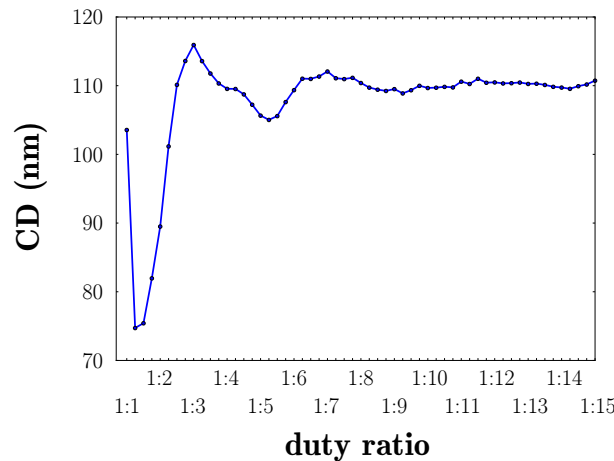


Figure 2.19: Iso/dense bias for line/space pattern. The duty ratio denotes the ratio between the feature size (line width) and the space.

The iso/dense bias can lead to a vast intra-die variation of the CD. Therefore it has to be taken into account and counteracted when designing the photomask.

- *CD non-linearity*: As discussed in Section 2.2.5, mask features close to the resolution limit show a non-linear relation to their images on the wafer. This has not only to be taken into account for

the CD control, it also requires a more rigorous mask error treatment of features with a MEEF substantially greater than one.

- *Line shortening and line-end rounding:* The loss of high diffraction orders also leads to a rounding of the line-ends and to line shortening effects. That is, although the width of a line may be on target, its length is not. Since the impact of a diffraction order loss increases with smaller features, even more rigorous countermeasures have to be taken in a low k_1 scenario. Moreover, OPC corrections targeted at line shortening have to maintain the width of the feature. Figure 2.20 demonstrates line shortening. The target line width is 90 nm which is clearly met. However, the designated length of 500 nm is not reached.

This effect is critical as it degrades the overlay accuracy, leading, for example, to a misplacement of interconnects.




Figure 2.20: Line shortening due to line-end rounding. (a) Although the CD is on target, the length of the line is not. Shortening can be corrected by (b) serifs or (c) hammer head assist features.

- *Corner rounding:* Similarly to line-end rounding, the filtering of diffraction orders also leads to noise at the corners of features. For example, contacts and contact holes are typically realized as rectangles on the photomask. Due to the corner rounding, the image, however, resembles a circle. Especially elbow structures as depicted in Figure 2.21 are highly sensitive to corner rounding effects.

The term OPC denotes photomask-design-related countermeasures to reduce optical proximity effects. Depending on the proximity effect to suppress, different techniques can be employed:

The two first proximity effects (*iso/dense bias* and *CD non-linearity*) impact the linewidth. They can hence be considered one-dimensional effects. An obvious correction is to bias the linewidth according to their context, for example, the pitch (*iso/dense bias*) or the absolute CD (*non-linearity*). Linewidth adjustment on the mask can only be performed within constraints presented by the design and the resolution of the mask writer. An alternative technique, termed half-tone biasing, has been proposed by Nakagawa et al. (1999).¹¹ There the edge of a line is subdivided into segments with a period below the resolution limit of the projector. The elements are not resolved individually but an average intensity (and phase) is transmitted. Tuning the size and the period of the edge elements allows for a linewidth control below the resolution of the mask writer. Especially when aiming at equalizing the size of dense and isolated features, linewidth biasing is not ideal, as it is effective only for a narrow focal range. In order to maximize the common process window of dense and isolated features, another technique, using

¹¹(also see Wong, 2001, chap. 4, pp. 95–96).




Figure 2.21: Corner rounding at an elbow structure: (a) With the original mask, both the convex and the concave corner show a pronounced distortion. (b) Applying a serif at the lower left convex and a “mouse-bite” (anti-serif) at the opposite concave corner, significantly improves the image.

sub-resolution assist features (SRAFs) (also called scattering bars (SBs)), can be employed (Prouty and Neureuther, 2004; Garofalo et al., 1993; Chen et al., 1997). The basic idea is to condition isolated features such that they behave similarly to dense ones. This is achieved by adding elements that do not actually print but that lead to a diffraction scattering behavior comparable with dense features. Figure 2.22 demonstrates the application of SRAFs in order to improve the process window overlap of dense and isolated features. There an array of scatter bars is placed to both sides of an isolated line, leading to an improvement of its DOF and the exposure (threshold) latitude. The latter can be attributed to the fact that off-axis illumination, which was also assumed in this example, strongly increases the dose latitude of dense features but only marginally affects isolated ones. By transforming the isolated element into a quasi-dense structure, it can profit from the same RET measures as dense features.

Finding ideal assist set-ups for multiple pitches is a strong optimization problem, especially when combined with other RETs such as off-axis illumination and phase-shifting masks (Mansfield et al., 2000; Smith, 2002).¹² It is further complicated by the impact SRAF OPC has on design and manufacturability (Schellenberg et al., 2001; Liebmann, 2003).

The two remaining optical proximity effects, line shortening and corner rounding are strictly two-dimensional. As shown in Figure 2.20, line shortening can be compensated by introducing assist elements such as (b) serifs (Saleh and Sayegh, 1981; Fu et al., 1991) or (c) hammer heads, a specific form of serifs. If design and manufacturing permit it, a simple lengthening of the lines is another viable option to correct shortening effects (Wong, 2001, chap. 4, p. 100). Moreover, also off-axis illumination reduces optical proximity effects such as line shortening (ibid.). Serifs are also suited to restore straighter corners that are subject to corner rounding, as shown in Figure 2.21(b). An anti-serif, sometimes called “mouse bite”, is used in the opposite concave corner (c).

OPC is a key step when designing the chip as it is to appear on the photomask. Traditionally, it was performed on otherwise completed design data that have been generated under various non-lithographic aspects such as the electrical behavior, overlay and other manufacturability criteria. Today, lithographic process constraints are considered in earlier design stages, leading to “litho-friendlier” designs. Both scenarios prompt for computer-aided approaches that (semi-) automatically incorporate the lithography-specific corrections into the layout. They can be coarsely subdivided into two categories: rule-based and model-based OPC.

A *rule-based* approach makes use of look-up tables whose keys are design situations and whose entries are the due OPC modifications (Otto et al., 1994; Chen et al., 1997; Park et al., 2000). For example,

¹²An optimization routine for this problem will be discussed in Chapter 7.

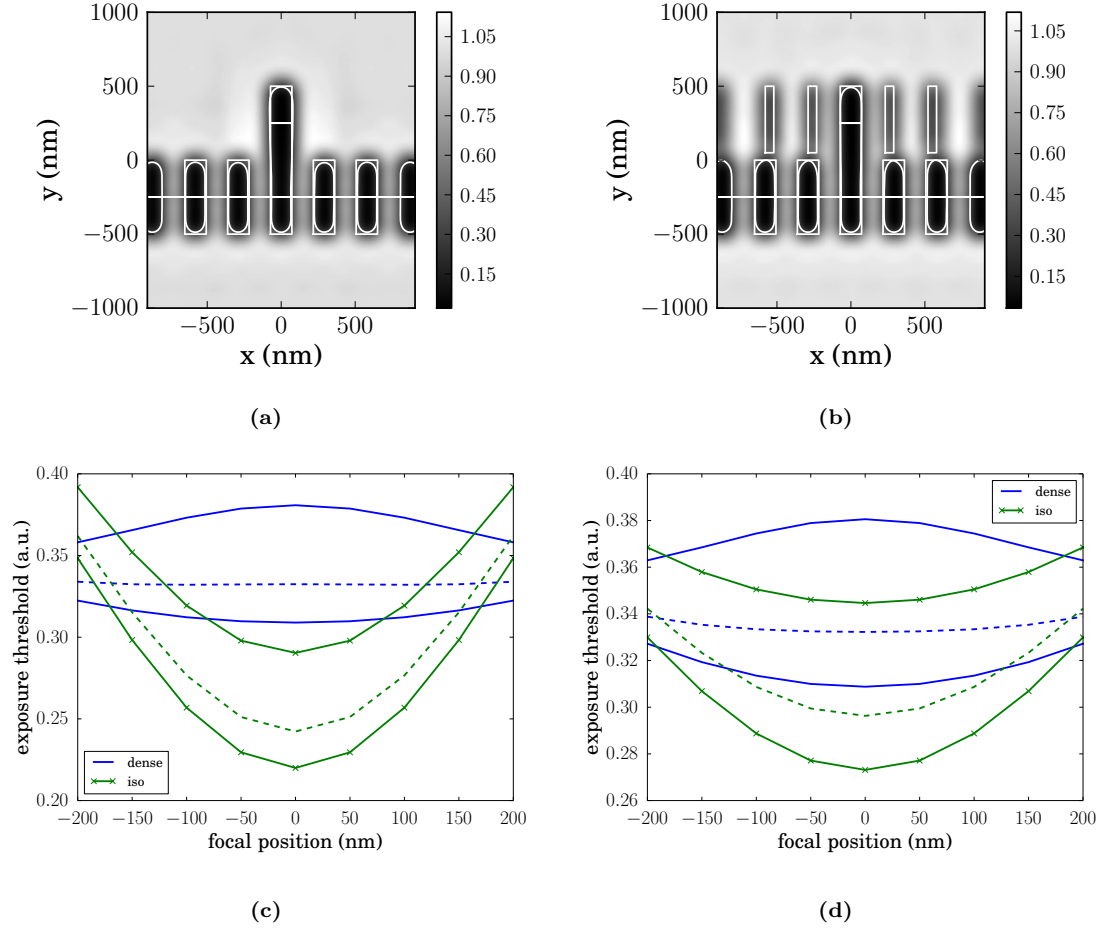


Figure 2.22: Use of SRAFs to maximize common process windows of isolated and dense features: Here a mixed isolated/dense lines/space pattern is imaged (a) without and with (b) SRAFs. The CD extraction position is indicated by the white horizontal lines. In the uncorrected case, no feasible common process windows for the upper isolated and the lower dense situation can be achieved (c). The addition of scatter bars leads to a harmonization of the process windows and thus to a larger overlap (d). An optimization procedure for the SRAF sizing and placement would be required to further increase the PW overlap.

A 193-nm wavelength, an NA of 0.8 and an annular illumination (inner radius: 0.3, outer radius: 0.5) were used in this example. A threshold of 0.3 was applied.

a table aiming at iso/dense bias corrections could contain as keys a number of edge-to-edge distances and linewidths, and as entries the edge displacement values that bring the respective CD on target (see Wong, 2001, chap. 4, p. 102). Such rules are, of course, strictly empirical. They have been found by experiment, simulation or simply by experience. Such a procedure is suited for situations in which only a limited number of simple corrections are required. Today, with an ubiquitous use of complex OPC schemes, a look-up table would require millions of entries, which not only have to capture the original proximity issues but also secondary effects caused by the OPC measures themselves.¹³ Especially the generation of the tables then poses an efficiency problem. Moreover, the tables are valid only under the same process conditions. Any change of illumination, NA, resist etc. immediately entails the compilation of a new set of rules.

Many OPC engines therefore use a more flexible *model-based* approach (Stirniman and Rieger, 1994; Cobb et al., 1996; Lucas et al., 2001; Kamat et al., 2002; Chen et al., 2007), consisting of a fast lithography simulator,¹⁴ capturing both optical and photochemical aspects (Patterson et al., 2005; Dolainsky et al., 1997)—often also a compact etch simulator—and algorithms for the layout manipulation. In most cases, model-based OPC is an iterative procedure in which an initial or intermediate layout is modified, aerial images and resists computations are performed, and the result is compared with a design intent. Schemes based on a Monte-Carlo approach (Kim et al., 2002), artificial neural networks (Huang et al., 2006), genetic algorithms (Li et al., 2008) and linear regression (Gu and Zakhori, 2008) were proposed over the past years. Figures of merit may include a multitude of local and global criteria for the quality of both the image and the layout, as well as manufacturing considerations (Kahng and Pati, 1999; Cobb and Granik, 2002). The iteration loop terminates as soon as a predefined requirements are met.¹⁵

Despite the improved flexibility, model-based approaches exhibit a number of issues that can be moderate by combining them with design rules. On the one hand, the iterative modification/verification loop leads to a time-expensive procedure, which can be shortened by limiting the degree of freedom, for example, by introducing constraints in form of rules. On the other hand, the placement and sizing of sub-resolution elements such as scattering bars or serifs can be done in more or less straightforward and rigorous manner using models. But an algorithmic approach to the generation of such patterns “from scratch” can hardly be done without rules (Chen et al., 1997; Liebmann et al., 2001; Brist and Torres, 2003; Mukherjee et al., 2005). In that sense many OPC tools follow a hybrid approach.

2.6.2 Off-Axis Illumination

As discussed before, the extension of the illumination source determines the partial coherence or filling factor which in turn constitutes an effective way of influencing the image behavior. Large partial coherence factors increase the number of orders collectible by the pupil (cf. Figure 2.11(b)). They are hence beneficial for patterns—such as lines/spaces—with a high spatial frequency. On the other hand, for isolated features a contrary situation occurs: (1) The large inclination angles at the outer radii of the source may entail a truncation of higher diffraction orders that also carry higher intensity magnitudes. The so resulted imbalance can lead to a degradation of the exposure latitude. (2) In addition, different angles of incidence result different through-focus behaviors of the respective images. The larger the range of incidence angles, the larger the variation of the out-of-focus images. Consequently, in the case of large partial coherence factors, the superposed image shows a significantly higher degree of blur than for smaller filling factors.

In conclusion, the selection of an adequate filling factor strongly depends on the feature. Wong (2001, pp. 78,79), for example, suggests a selection procedure based on the modulation transfer function (cf. Section 2.2.6): Given features with a principal normalized pitch \hat{p} ,¹⁶ find the filling factor that maximizes

¹³An algorithm based on such a procedure would indeed have to be iterative, with an unpredictable convergence behavior.

¹⁴Compare Section 3.2.6.

¹⁵This procedure closely resembles the mask optimization portion in source/mask optimization approaches, which will be discussed in Chapter 7.

¹⁶

$$\hat{p} := p \frac{\text{NA}}{\lambda},$$

where p is the pitch.

the contrast at frequency $1/\hat{p}$. The inflection point between high and low partial coherence factors hence lies at a spatial frequency of one. However, only in rare scenarios, extreme filling factors are suitable: (1) Often, IC layers consist of a multitude of different pattern frequencies, prompting for a compromise between large and small σ values, (2) practical considerations, such as robustness against aberrations and intensity fluctuations, or throughput maximization, limit the minimum partial coherence factor to values between 0.2 and 0.3 (Wong, 2001, p. 79), (3) additional resolution enhancement techniques such as OPC or phase-shifting masks strongly interdependent with the employed source.

Although large filling factors may under certain conditions improve the contrast of dense lines/spaces images, much of the effective source contributes only to their background intensity. The following example (adapted from Wong, 2001, p. 80) demonstrates this property: The target is to image a line/space pattern with equal line and space widths and a pitch of $5/8$ (of λ/NA), as shown in Figure 2.23(a). Using a large partial coherence factor of one, Figure 2.23(b) illustrates the small portions of the source that are actually conducive to two diffraction orders and thus to the pattern formation. The remaining areas of the source solely increase the background intensity, impairing the contrast of the resulted image. As a counteraction, one can employ sources that omit regions where no interaction between diffraction orders occurs. Two examples for such types of illumination sources are depicted in Figure 2.23(c). On the left side of the source, the left pole of a di-pole set-up is demonstrated. The geometry of the pole is determined such that it approximates the region of diffraction order interaction in the source: the inner radius (σ_i) is 0.6 and the aperture angle (α) is 60° . The right side shows the right part of a free-form source that exactly matches the interaction area. The bright dots indicate the source point sampling used for the following image computation. The performance of the different source types is evaluated by conducting an aerial image computation. Figure 2.24 shows the result for five different sources: (1) A circular illuminator with a partial coherence factor of 0.6, (2) a circular source with a σ of 1.0, (3) the previously motivated di-pole source, (4) the free-form source matched to the interaction region of the diffraction orders, and (5) the same free-form source with y polarization. In addition to the best focus, images for a 400 nm defocus are presented. With the circular illuminator and a partial coherence factor of 0.6, the pattern collapses completely since no appreciable interaction between diffraction orders occurs. A sigma of 1.0 slightly improves the situation. However, due to the large intensity offset resulted by the virtually superfluous regions of the source, a stable pattern is still not achievable. The pattern is drastically enhanced by the introduction of a di-pole illuminator which at best focus shows the same good performance as the unpolarized free-form source. Only at a large focus off-set of 400 nm the free-form source significantly outperforms the di-pole in terms of contrast. The ultimately best performance in this comparison is given by the y -polarized free-form source, for which a very stable pattern is obtained at best focus, and which still displays a low but distinct contrast under the large defocus. The previous example demonstrates the selection of off-axis illumination sources for one-dimensional line/space patterns. For patterns such as contacts that are intrinsically two-dimensional, similar considerations lead to quadrupoles as an ideal source geometry. Of course, in reality different pattern types have to be imaged at the same time. Therefore, a compromise between different illumination set-ups has to be found. In addition, multi-pole illuminations are radially asymmetric and may so aggravate other resolution enhancement efforts (Kamon et al., 1991; Wong, 2001). For these situations, an annular illumination source, in which a disk in the center of a circular illumination is omitted, restoring radial symmetry but also re-introducing a level of background intensity. Typical illumination set-ups and their parameters are summarized in Figure 2.25.

Another important parameter of the illumination is its polarization. Especially in the regime of large NAs, the polarization state significantly impacts the image performance. For the discussion (and for modeling purposes, as will be shown in Section 3.2.2), it is convenient to decompose vibration directions of an electric field into two modes:

- (1) The *transverse electric* (*TE*) component denotes the vibration direction perpendicular to the meridional plane (or plane of incidence) that is spanned by the propagation vector (\mathbf{k}) and the optical axis (z):
- (2) The *transverse magnetic* (*TM*) component describes vibrations that coincide with the meridional plane and that are orthogonal to the TE direction.

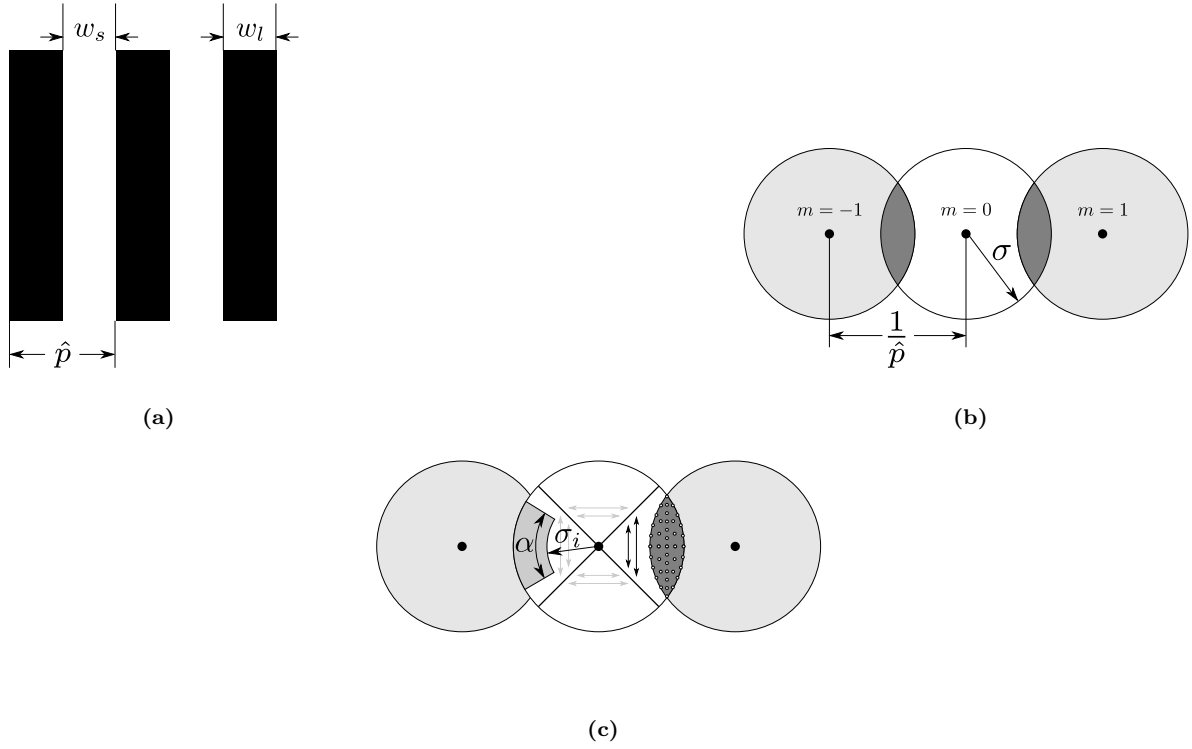


Figure 2.23: Comparison of different illumination settings for a lines/spaces pattern. (a) The pitch of the mask (\hat{p}) is $\frac{5}{8} \frac{\lambda}{\text{NA}}$, having equal line (w_l) and space widths (w_s). (b) The zeroth diffraction order (m) lies in the center of the extended source, from which the $\pm 1st$ orders have a distances of $\frac{1}{\hat{p}}$. The dark-gray-shaded areas denote the regions of the source that contribute to two diffraction orders for a filling factor (σ) of one. In (c) two examples of sources that omit regions not contributing to the pattern formation are demonstrated, depicting the left part of a dipole approximately encompassing the area of diffraction order interaction and the right part of a free-form source that exactly matches the interaction region—where the bright dots indicate the source point sampling used for the image computations below. The arrows symbolize optional polarization states of the source that will be explained below.

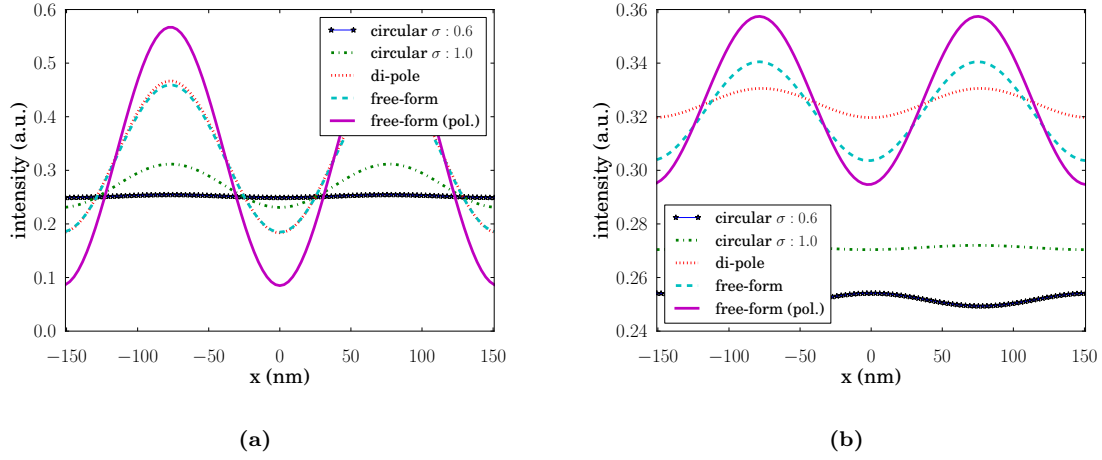


Figure 2.24: Intensity distribution under different illumination conditions (a) in-focus and (b) for a 400 nm defocus. The di-pole and the circular configurations are unpolarized, the free-form source has been evaluated for an unpolarized and an y -polarized state. The circular illumination scheme only provide a marginal contrast, already when evaluated in the focal plane. Without defocus, the unpolarized free-form and the di-pole show a similar intensity distribution. But under a 400 nm defocus, the free-form source exhibits a significant intensity advantage over the dipole. When applying y -polarization, the free-form clearly outperforms all other (unpolarized) configurations.




Figure 2.25: Parametric representation of typical illumination types: (a) Circular illumination, defined by its radius σ . (b) Annular illuminator, defined by outer and inner radii σ_o and σ_i . (c) Traditional multi-pole set-up, where parameters are the number of poles, the radius of the circle on which the centers of poles sit (σ_o), the blade angle by which this circle is rotated (θ), and the radius σ_p of the individual poles. (d) State-of-the-art illuminators are commonly implemented using an aperture approach. There the poles are realized as segments of a circle, defined by the inner and outer radii (σ_i, σ_o) of the segment, its aperture angle α and the blade angle θ describing the rotation of the poles.

The TE and TM contributions are also called *s-polarization* (from the German word *senkrecht*: perpendicular) and *p-polarized* (parallel), respectively.




Figure 2.26: Two-beam imaging. While the *p* components strongly depend on the angle of incidence no such dependence occurs for the *s* components.

To illustrate the impact of polarization, we shall consider the two-beam imaging case that is illustrated in Figure 2.26: The TE (or *s*) components, pointing out of the figure, overlap independently of the incidence angle. In contrast, the interaction of the TM (or *p*) components decreases with growing angles, leading to decrement of their contribution to the pattern formation and hence to a degradation of the image. The previous example (cf. Figure 2.24) clearly demonstrated that even for NAs smaller than one, polarized illumination sources are suited to improve the image contrast, even more when considering out-of-focus images. There, the source was TE polarized, which in the one-dimensional case, that is, without consideration of the *y*-extent of the line/space pattern, is equal to a *y*-polarization state. The different polarization states are depicted in Figure 2.27. In addition, sub-figure (e) shows a segmented illumination source, in which the upper and lower partitions are *x*-polarized, and the left and right partitions are *y*-polarized. This set-up approximates the TE condition and is often employed because of technical feasibility (Jasper et al., 2006).




Figure 2.27: Common illumination polarization states.

Although the particular benefit of the polarization state of the source has to be evaluated in conjunction with both the photomask and the projection lens system—both significantly influence polarization—, it is generally assumed that TE or *x* and *y* approximations thereof can be considered the beneficial polarization state. Figure 2.28 shows typical illumination geometries and corresponding polarization set-ups. Since the past recent years, off-axis illumination sources are commonly realized using diffractive optical elements (DOEs). DOEs are computer-generated holograms that re-shape the input light to form a desired intensity distribution (Himel et al., 2001). This is achieved by a grid structure consisting

of pixels with different phase levels. DOEs can be manufactured either in a direct-write process or, suited for high-volume production, in a modified projection lithography process (Leonard et al., 2008). The phase levels are realized by employing a sequence of exposure, develop and etch steps. Both the resolution of the grid and the number of phase levels strongly depend on the requirements of the application. Typical phase levels are 2, 4 or 8, and a grid resolution of 200 times 200 is common (Carriere and Hudson, 2009). Due to their light redirecting, DOEs exhibit a much improved light efficiency of up to 94% for the eight level design compared with formerly utilized aperture stops. Although this property increases the achievable throughput, it drastically elevates the risk of intensity peaks and consequently of damages of the condenser lens. In addition to the parameterized sources such as annular or multi-pole illuminators, DOEs allow for custom illumination set-ups, often termed free-form sources (see Figure 2.23(c) of the previous example). Due to the considerable gain in flexibility associated with DOEs, this technology can be considered a key enabler of source/mask optimization.

The main disadvantage of this approach is that each custom source requires the fabrication of a tailored DOE. In order to change the illuminator, the DOE has to be replaced and the condenser lens set-up has to be adjusted. An even more flexible approach, using an array of programmable micro mirrors, was recently proposed by (Mulder et al., 2010). The shape of the illuminator is determined by the tilt of the micro mirrors, whose number is in the order of several thousands. In order to guarantee the accordance of the actual state of the source with the target source, an *in situ* measurement and control mechanism is incorporated into the proposed approach. Switching of the illumination shape involves a control loop in which the mirrors are adjusted until the given target shape is achieved. Similarly to DOEs, there is no significant light loss: For each source shape, all mirrors, which have a coating that is highly reflective at a wavelength of 193 nm, are utilized. According to the authors, preliminary tests proved both the short-term reproducibility of the source shape and the long-term durability of the mirrors and the mechanical components.




Figure 2.28: Chart of typical illumination set-ups (taken from Jasper et al., 2006).

2.6.3 Phase-Shifting Mask

As discussed before, different mask patterns interact through interference in the lithographic image. The result of this interaction is determined by the incidence angle and the form of the individual wave trains. One way to improve resolution is to bias the phase of these wave trains. This can be done by utilizing phase-shifting mask elements. The three commonly used techniques are: *alternating phase-shifting masks (AltPSM)*, *attenuated phase-shifting masks (AttPSM)* and chromeless masks that are used in *chromeless phase lithography (CPL)*.

Alternating phase-shifting mask

In alternating phase-shifting masks (AltPSMs), a phase shift between two bright areas encompassing a dark area is applied (Hänsel and Polack, 1976; Levenson et al., 1982; Shibuya, 1982; Wong, 2001). In the schematic shown in Figure 2.29, a binary chromium-on-glass (COG) mask (a) is compared with an AltPSM (b). While in the COG case both bright regions show the same phase (0 rad.), the AltPSM incorporates a phase shift of π rad. for the right transmitting region.




Figure 2.29: Comparison of (a) chromium-on-glass mask (COG) and (b) alternating phase-shifting mask: While the bright features in the COG mask result the same phase (0 rad.), in the AltPSM, the right bright feature leads to a phase shift of π rad.

The phase shift leads to a destructive interference of light transmitted through the bright regions, resulting in an intensity decline in the intended dark area in-between the clear regions. In other words, the light-blocking property of the opaque feature is amplified by the destructive phase interaction of the surrounding transparent regions. Figure 2.30 shows the intensity distribution under an AltPSM in comparison with a COG mask: a 97-nm line and space pattern is imaged (λ : 193 nm, NA: 0.8, σ : 0.5). The contrast is significantly improved.

One of the reasons for the improved imaging behavior of AltPSMs lies in the fact that there is no transmission of the zero diffraction order. While—without off-axis illumination—COGs require three orders, the zeroth, and the plus and minus first orders, to be collected for a pattern to appear in the image plane, AltPSMs depend only on the two first orders, thus in an ideal case, allowing for a perfect two-beam imaging scenario. Figure 2.31 illustrates this property in the spatial frequency domain. It shows that in the AltPSM case not only the zeroth order but all even orders vanish. As in the COG case, the distance between the propagating orders is one divided by the normalized pitch (p). However, the frequency shift results in a distance of $1/p$ between the mandatory plus/minus first orders as opposed to $2/p$ for COGs.

This is often referred to as frequency-doubling as it significantly increases the achievable process window (Figure 2.32) and consequently allows for denser patterns to be printed.

Another advantage of AltPSMs is the decreased mask error enhancement factor (MEEF; see Section 2.2.5). The destructive interference between the transparent regions render this mask type less sensitive to small size deviations in the opaque regions. Although less sensitive to mask errors, however, AltPSM complicate the detection and repair of mask defects.

The AltPSM manufacturing process is significantly more complicated than the production of COG masks as it involves an additional alignment and etch step in order to generate the phase shift. As shown in Figure 2.33, after the chromium absorber has been deposited, like in the COG manufacture process, an additional etch step is required, in which a trench in the silica substrate is generated. This trench




Figure 2.30: Intensity distribution under COG and AltPSM. The aerial image presented here was computed using a thin-mask approach, which does not account for rigorous effects such as intensity imbalances, discussed later.

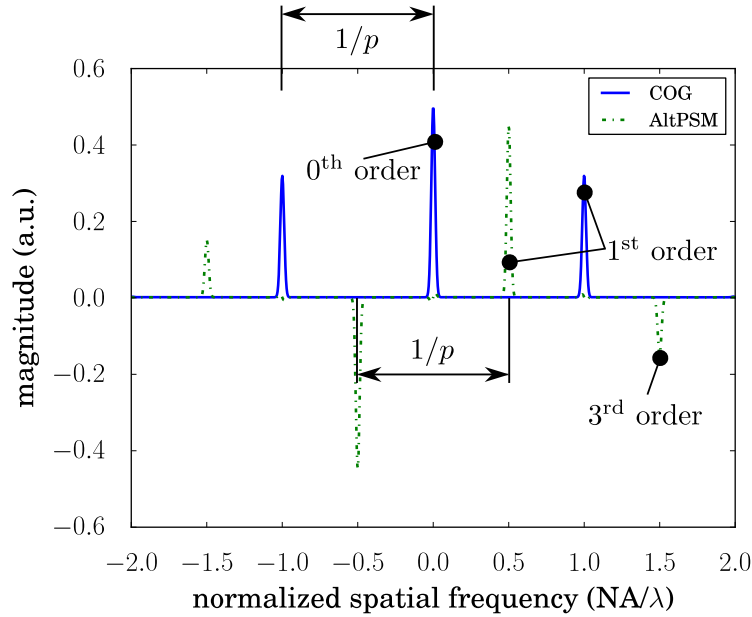


Figure 2.31: Comparison of the amplitude distribution in the spatial frequency domain of a COG and an AltPSM. While the COG requires three orders to be collected by the pupil, AltPSMs do not exhibit even orders. Thus, in an ideal scenario, a two-beam imaging can be achieved. Moreover, the plus-minus first orders of the AltPSM have only half the distance to the DC order, compared with the COG—a phenomenon often referred to as frequency doubling.




Figure 2.32: Process Window for COG and AltPSM. The AltPSM shows a significantly improved exposure latitude.

generates the phase shift. For example, for a phase shift of π , a path difference of half the wavelength has to be attained. Thus,

$$\begin{aligned}
 n_{1,\lambda}d_e - n_{2,\lambda}d_e &\stackrel{!}{=} \frac{\lambda}{2} \\
 \Rightarrow d_e &= \frac{\lambda}{2(n_{1,\lambda} - n_{2,\lambda})} \\
 \begin{matrix} n_{1,\lambda} \approx 1.563 \\ n_{2,\lambda} \approx 1 \end{matrix} &\Rightarrow d_e \approx \frac{\lambda}{1.126},
 \end{aligned} \tag{2.15}$$

where $n_{1,\lambda}$ is the wavelength-dependent index of refraction of the substrate and $n_{2,\lambda}$ is the refractive index of the medium that surrounds the mask. For fused silica and a wavelength of 193 nm, the former has a value of about 1.563. The refractive index of air, the medium that comprises the mask, is about 1, yielding an etch depth (d_e) of about 171.4 nm.




Figure 2.33: Generation of phase shift in AltPSM mask. After a manufacture of a chromium-on-glass.

Despite the improved imaging behavior due to an improved depth of focus and exposure latitude, AltPSMs impose a number design problems. One of them is an intrinsic difference between the intensity resulted from the phase-shifted openings. This property can be numerically verified when applying an rigorous electro-magnetic field (EMF) mask near-field simulation to the AltPSM. Reflections from the absorber and from the sidewalls of the substrate lead to a significant amount of standing waves. The impact of these standing waves depends on the opening depth of the features. In fact, it can be shown that the deeper phase-shifting region more severely suffers from destructive interference due to back-reflection (Wong and Neureuther, 1993). The intensity underneath this region is consequently lower than

the one of the nonshifted feature. This intensity imbalance leads to a CD difference between phase-shifted and nonshifted (bright) spaces, and more critically, to a placement error of the (dark) lines. Moreover, diffraction from the substrate and absorber edges may cause that the desired phase-shift of π is achieved only in the nominal focal plane but is different from π for out-of-focus positions. Various countermeasures to reduce intensity imbalances have been proposed, such as increasing the width of the phase-shift region (Wong and Neureuther, 1994), employing deep trenches, producing a π and a 2π phase-shift (Tanaka et al., 1996; Gordon et al., 1998), isotropic etching to obtain an underetched phase-shift region (Pierrat et al., 1992), and combinations of size biasing and isotropic etch steps (Ferguson et al., 1998).

Another problem is, coarsely speaking, to separate phase regions from each other in order to guarantee the manufacturability of the mask but also to avoid phase conflicts and resulted artifacts. Figure 2.34 demonstrates a situation in which an image of a bright-field AltPSM would show unintended intensity minima that are due to phase conflicts.




Figure 2.34: Phase conflicts with AltPSMs. In this bright-field lithography example, a phase shift is applied to improve the imaging behavior of for a line pattern. In addition to the intended destructive interferences in the image of the dark feature, low intensity artifacts occur wherever the background and shifted phase clash.

A common approach to cope with phase-conflict-induced issues is to apply a second exposure step. In the previous example, a second bright-field exposure using a binary mask could be applied to retroactively brighten up the unwanted dark regions. Such a procedure, however, dearly buys the artifact removal—at the price of a significant contrast degradation due to flare. Performing the trim exposure with a complementary COG mask—in the previous example, that is a dark-field COG mask, transmissive only in the regions where phase-shift errors are to be removed—drastically lowers the insertion of flare. The example in Figure 2.35 illustrates how this principal is pursued in state-of-the-art AltPSM processes (Rieger et al., 2001): The intended design is split into a less critical trim pattern and an AltPSM that is to generate the critical line widths. While the trim exposure is done with a bright-field COG mask, yielding an unexposed envelope of the pattern, the AltPSM step for the critical lines is conducted with a dark-field mask.

In addition to comparable double exposure techniques using trim masks, which have been proposed, for example, by Suzuki et al. (1999); Fritze et al. (2001); Levenson et al. (2001), phase conflicts can be reduced by incorporating softer phase shifts. Ohtsuka et al. (1992) have presented a technique, called *conjugate twin shifter* with three phase values: 0, $\frac{1}{2}\pi$ and $1\frac{1}{2}\pi$. The destructive interference for the dark regions is achieved by the $\frac{1}{2}\pi, 1\frac{1}{2}\pi$ phase transition. Under well-behaved conditions, the phase shift from 0 to both $\frac{1}{2}\pi$ and $1\frac{1}{2}\pi$ still generates a high enough intensity in the image. Therefore the

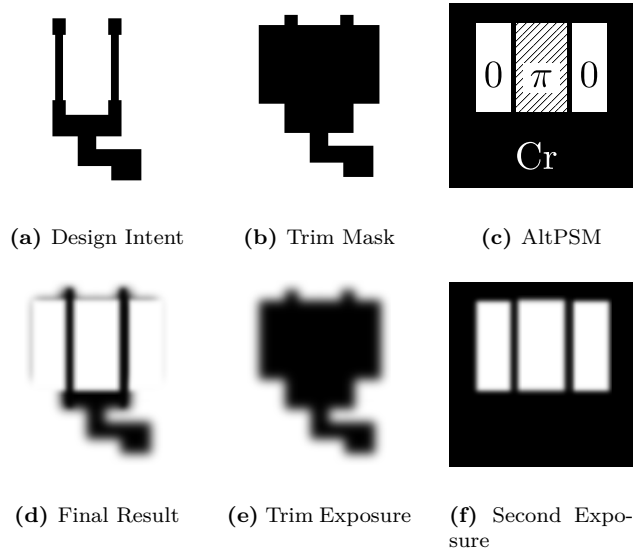


Figure 2.35: Double exposure for AltPSMs: (a) Target pattern; (b) a bright-field trim mask is used to produce the less critical envelope of the target pattern (e); (c) a dark-field phase-shifting mask is applied in a second exposure (f) to refine the pattern, i.e., to generate the critical aspects of the feature, yielding the final result (d). Example taken from Rieger et al. (2001).

mask can be designed such that no trim exposure is required. A side-effect of this technique, however, is the introduction of focus-dependent asymmetries (Ohtsuka et al., 1992; Wong, 2001), which in the extreme case may entirely foil the resolution enhancement purpose of AltPSMs.

A slightly different approach was proposed by Pfau et al. (1992). There the π -phase-shift region is surrounded by a gradual phase transition from π to 0 at edges at which no destructive phase interaction is intended. Nistler et al. (1992) aim at achieving a similarly smooth phase transition by using a discrete number of phase region at the edges, each of which has a fixed phase value. In their original work, two such phase levels are proposed: $2/3\pi$ and $1/3\pi$.

Yet another technique—called *self-aligned phase shifters* (and later rim PSMs)—to eliminate the design issues associated with alternating phase shifting masks was earlier introduced by Nitayama et al. (1989); Terasawa et al. (1989); Yanagishita et al. (1991). There the dark regions of a COG mask are surrounded by phase-shifting elements. These elements lead to an edge contrast enhancement through destructive interference. Although derived from AltPSMs, rim PSMs are referred to as weak phase-shifting masks, since the mask design does not depend on the phase difference of adjacent bright features; rather the phase-shift is employed to amplify the intensity blocking intended by the opaque mask regions. Rim PSMs have been shown to be compatible with other resolution enhancement techniques such as off-axis illumination (Brunner, 1993). Mainly because of manufacture reasons and because of the popularity of another type of weak PSM, the attenuated phase-shifting mask—introduced in the next section—rim masks do not currently play a prominent role in lithography (Lin, 1992b).

Another issue arising from the utilization of AltPSMs is their sensitivity to projection lens aberrations (Wong et al., 2001; Slonaker, 2001; Mak et al., 2004). Therefore an even stricter aberration control of the projection system, especially in the high-NA regime, is critical for the successful application of AltPSMs. In addition, similar countermeasures as for the intensity imbalance issue can be applied to reduce potential placement errors (e.g., see Mak et al., 2005). There is also a strong sensitivity of AltPSMs on polarization (Flagello et al., 2004; Erdmann, 2005; Bubke et al., 2006; Gleason and Cheng, 2006). Whether and how polarized light can be applied in conjunction with AltPSMs is still subject to research. Bubke et al. (2006) conclude that no significant benefit is provided by employing polarized

light. Moreover, additional polarization-dependent balancing techniques have to be compassed (this observation was confirmed, for example, by Gleason and Cheng, 2006).

Attenuated phase-shifting mask

Like the previous mask type, also attenuated phase-shifting masks (AttPSMs) utilize a destructive interference presented by a phase shift of adjacent elements in order to enhance contrast (Terasawa et al., 1991; Lin, 1992a; Wong, 2001). Neglecting other RET measures, the efficiency of the zeroth diffraction order—which, as shown before, does not contribute to the image modulation but rather leads to a contrast decline—is reduced by allowing a certain portion of light transmittance through the dark regions. These regions implement a 180° -phase-shift that leads to the desired destructive interaction. Since the amount of light through the dark areas is only a fraction of the fully transparent regions, the coincidence of the phase-different waves does not result in a complete cancellation but an intensity reduction.¹⁷

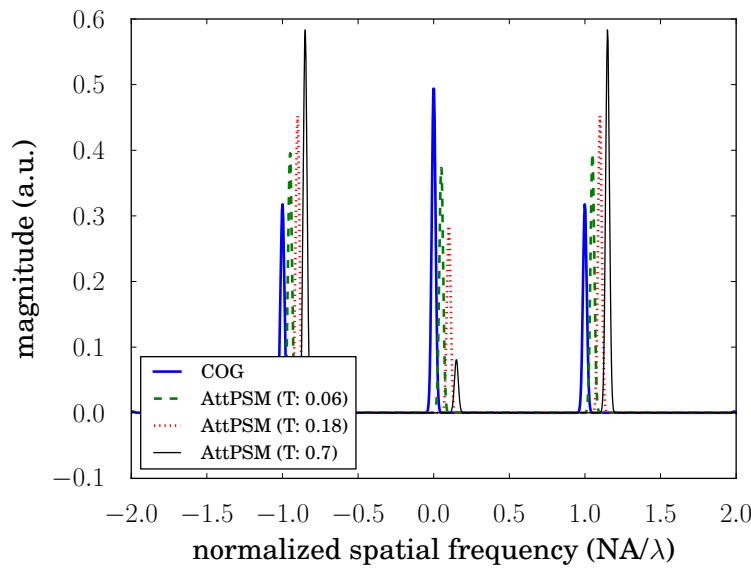


Figure 2.36: Comparison of the diffraction order efficiencies of a COG and AttPSMs (1:1 duty ratio): In contrast to the alternating PSMs, both mask types exhibit the same order distribution.^a Depending on the background transmission value (0.06, 0.18 and 0.7), the zeroth order can be attenuated while at the same time the higher diffraction orders are accentuate.

^aFor visibility's sake, the individual graphs are plotted slightly shifted in x direction as they actually overlap each other.

Figure 2.36 demonstrates that AttPSMs qualitatively show the same frequency response as binary masks. However, while for conventional masks the maximum magnitude lies in the zeroth order, the background transmission of AttPSMs directs more energy into higher diffraction orders. At a transmission as low as 4 to 6 percent relative to the transmission through transparent regions, the zeroth and the first diffraction orders are already balanced. Higher values are suited to further diminish the zeroth order but may also lead to an unintentional print of dark regions. This type of artifacts is commonly called sidelobe printing. Background transmission values of 18–19 percent (high transmission AttPSMs)

¹⁷Attenuated phase-shifting masks are hence often termed weak PSMs, in contrast to alternating PSMs, which are consequently referred to as strong PSMs.

constitute the absolute maximum for AttPSMs. The fictive value of 0.7 (70 percent) presented in Figure 2.36 would require additional exposure steps, similarly to AltPSMs. We shall see later (Section 2.6.3) that the CPL approach can be applied such that it can be regarded as an AttPSM with a background transmission of 1.




Figure 2.37: Aerial image of a 120-nm isolated bright feature, using a chromium-on-glass and an attenuated phase shifting mask with 6 percent transmission.

AttPSMs are well suited for imaging isolated bright features such as isolated lines or contacts. As shown in the example in Figure 2.37, in which the image of an isolated 120-nm bright feature under 193 nm illumination with a sigma of 0.3 and an NA of 0.8 is depicted, the edge between the dark area and the contact is sharpened by a pronounced intensity increase-decrease-increase succession. The steeper NILS, which it translates to, leads to a significant improvement of achievable process windows, as shown in Figure 2.38. However, the higher average intensity of the AttPSM and the brightened fringes surrounding the feature also increase the risk of sidelobe printing.

There are two types of attenuated phase-shifting masks (for a comparative study, see, for example, Yuana and Polb, 1995): (1) The *non-embedded AttPSM* (also *alternating aperture AttPSM*) in which the phase shift is realized through an etched trench in the fused silica substrate, whereas the transmission reduction of the dark regions is realized by a chromium layer and (2) the *embedded AttPSM*. Here both the phase-shift and the transmission reduction are implemented by a deposition of a film layer (for example, molybdenum silicide) with appropriate optical properties. Figure 2.39 demonstrates the different mask types, including a comparison with a conventional COG mask. The non-embedded AttPSM type presents numerous manufacture and inspection issues, mostly as it involves multiple etch steps: to thin the chromium absorber in order to adjust the transmission level, to obtain the transmissive apertures, and finally to apply the phase shift. While the implementation of the phase shift in the quartz substrate leads to similar issues as in alternating PSM case, including the requirement for a rigorous transmission and phase balance control, implementing a reproducible chromium thinning step—conducted mostly through wet etch—constitutes an even more severe technological challenge.

In contrast, the use of absorbers with embedded phase-shifting properties is suited to significantly reduce manufacture constraints. A wide range of absorber materials, have been proposed (Popa and Jonckheere, 1995; Jonckheere et al., 1994; Kalk et al., 1994; Ushioda et al., 1996; Matsuo et al., 1997;




Figure 2.38: Process window for COG and AttPSM for an isolated contact hole.

Reynolds et al., 1998; Lin and Loong, 1999; Loong and Lin, 2000).¹⁸ In addition to the appropriate optical material properties, specifically the π phase-shift and the target transmission, the processing properties amount to the proper material selection. Especially durability, film stress, substrate adhesion, inspectability and the existence of suited resists are important criteria (Smith et al., 1996; Carcia et al., 1996). Moreover, well-behaved etch characteristics and the availability of suited etch processes have to be guaranteed (Peng and William J. Adair, 1998; Choi et al., 2002). As an alternative to a single layer approach, Becker et al. (2003) proposed the use of an absorber stack consisting, for example, of a bilayer.




Figure 2.39: Comparison of (a) chromium-on-glass mask (COG) and (b) AttPSM in which the dark regions exhibit a phase shift of π and a low transmission, for example, of 7 percent. State-of-the-art AttPSMs typically use a molybdenum-silicide (often short MoSi) layer for the low transmission region.

Estroff et al. (2005) conducted a study on the polarization behavior of both COG and attenuated phase-shifting masks: While binary patterns act as TM polarizers only for deep sub-wavelength features, embedded AttPSMs can be said to filter out the TE polarized light regardless of pitch and feature sizes. This property leads to a degradation of the image contrast, which becomes even more severe in the regime of high projector NAs as employed in today's lithography. The authors conclude that any aspect of the mask, especially of the absorber, including thickness, material properties and side-wall angles, impact its polarization characteristics, the quest for a suited material should be carried on. Other comprehensive study on the impact of the optical material properties and the thickness of the absorber

¹⁸Some of the authors, e.g., Jonckheere et al. (1994), propose a hybrid approach. That is, the phase-shift is not completely embedded in the absorber but partly achieved by a quartz trench.

both on the imaging performance and on the MEEF under a high NA and polarized off-axis illumination was performed by Yamamoto et al. (2006) and Sato et al. (2007). In another study, Yoshizawa et al. (2006) showed that the TM polarizing effect of AttPSMs can be drastically reduced by employing a bilayer stack in which the transmission reduction is achieved by a metal compound and the phase shift is realized through a transparent material.

An investigation on the impact of different feature types on the diffraction efficiency behavior of AttPSMs was conducted by Smith (2002). Figure 2.40 shows the magnitudes of the primary diffraction orders under varying feature densities. In this example, the ratio of a line/space pattern is varied and the diffraction efficiency for the primary orders is illustrated. Here the duty ratio is defined as space width versus line width. That is, given a fixed pitch, the relative space width is varied from narrow (isolated space) to wide (isolated line). An increase of the space width also elevates the amount of energy directed into the zeroth diffraction order (Figure 2.40(a)). Moreover, due to a greater destructive phase interaction, higher background transmission values lead to a lower phase-normalized magnitude. The situation is reversed for the first diffraction order, whose magnitude increases under higher transmission (b). The diminishing impact of AttPSMs on the zeroth order and the increase of the first order magnitude with larger space widths, leads to a decreased interference contrast and consequently to an intensity off-set. The so raised risk of sidelobe printability as can be demonstrated by the following sidelobe intensity (I_s) estimation (Smith, 2002):

$$I_s := (A_{-1} + A_1 - A_0)^2,$$

where A denotes the magnitude of the respective diffraction order. Figure 2.40(c) depicts the corresponding graph. As shown, the sidelobe intensity is greatest at a semi-dense space duty ratio of about 1:2. Smith hence proposes the introduction of a forbidden pitch criterion for AttPSMs, stating that this mask type can be plainly applied only to certain feature layouts. Layouts with high sidelobe sensitivities require the application of additional RET or OPC techniques. Otherwise an AttPSM may even show a worse performance than COGs. Moreover, the study also illustrates that high transmissive AttPSMs bear a high risk of sidelobe printing for a wide range of duty ratios. State-of-the-art AttPSMs therefore utilize background transmission values of 6 or 7 percent. The problem of sidelobe artifacts may be aggravated in the presence of projector aberrations (Brunner, 1997; Otaka et al., 1997; Wong, 2001).

Because of their qualitative comparability with conventional binary masks, AttPSMs are compatible with the same type of resolution enhancement techniques such as OPC or off-axis illumination, which is in contrast to strong phase-shifting masks (Ronse et al., 1994; Ahn et al., 1995; Jonckheere et al., 1996; Wong, 2001). Of course, the sidelobe printing problem has to be taken into account when applying RET. Moreover, the order balancing property (as depicted in Figure 2.36) may provide an additional benefit for off-axis illumination. If, for example, the illumination set-up leads to a shift such that only the zeroth and the first order are collected, the difference of diffraction order magnitudes in COGs prohibits an ideal two-beam interference scenario. Attenuated phase-shifting masks, with an appropriate background transmission, can then be applied to balance the respective orders and to hence recover two-beam imaging.

As the introduction of EUV lithography is drastically delayed, sophisticated resolution enhancement techniques are required at a very early stage. In addition to off-axis illumination, attenuated phase-shifting masks (Sugawara et al., 2003; Chen et al., 2004; Shin et al., 2009) or hybrid masks (Kang and Hwangbo, 2007) are potential candidates. Like binary EUV masks, AttPSMs for EUV make use of multilayer mirrors in order to maintain the required reflectance levels for the bright regions (see Section 2.3). Thus fabrications demands are considerably higher compared with optical AttPSMs (Ghandehari et al., 2001; Jung et al., 2009).

Chromeless phase lithography

In contrast to the previously discussed phase-shifting mask types, CPL makes use of fully transmissive masks. To achieve low intensities, it therefore solely depends on the destructive phase interaction of phase-shifted regions on the mask (Toh et al., 1991b; Watanabe et al., 1990). This approach exhibits the

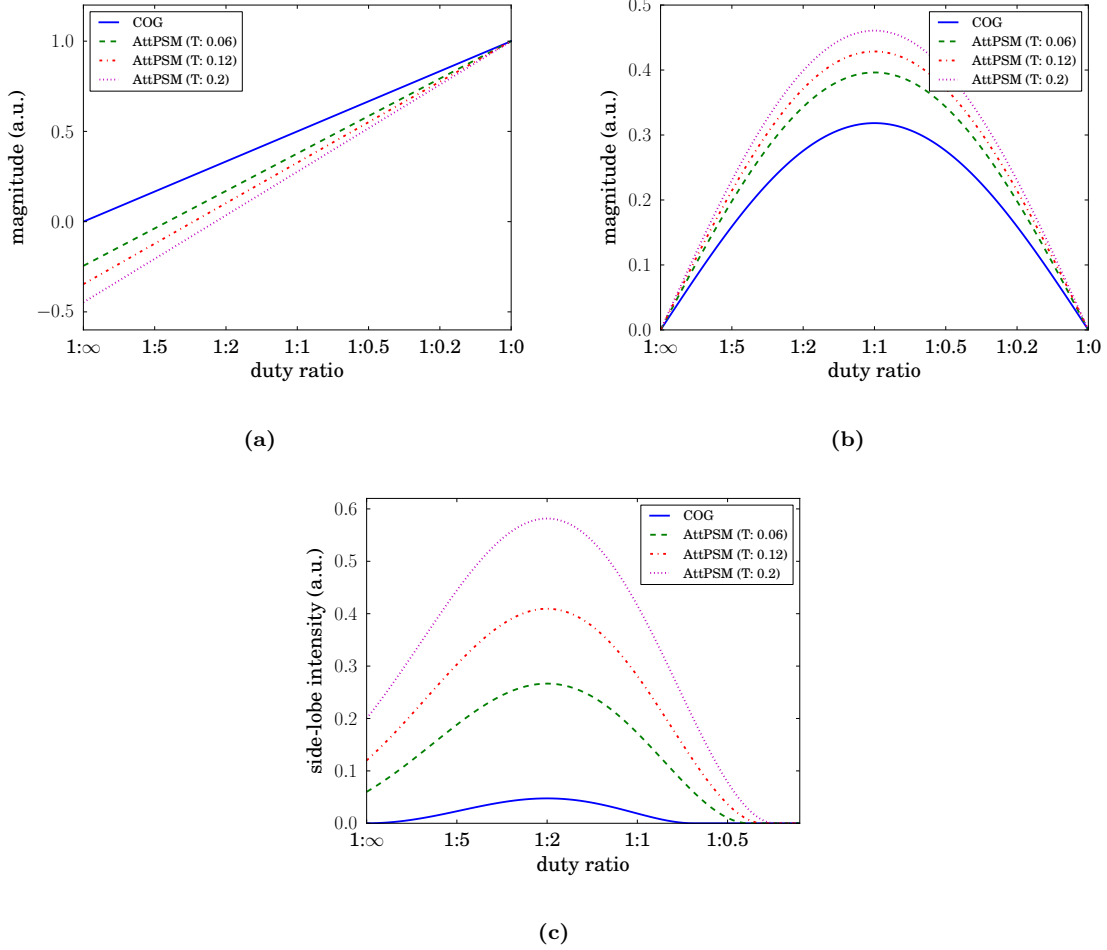


Figure 2.40: AttPSM diffraction efficiency of primary orders for different duty ratios and background transmission value (0.06, 0.12 and 0.2). (compare Smith, 2002). (a) Phase-normalized magnitude of zeroth order and (b) first order. (c) Sidelobe intensity for the same duty ratios. For all mask types, semi-dense spaces present the highest risk of sidelobes, rendering AttPSM better qualified for (semi-) isolated features. The high sidelobe intensity at a duty ratio of about 1:2 is due to a loss of interference contrast. The increased first order does not encounter a zeroth order with a corresponding magnitude, resulting an intensity offset and hence to artifact printability.

potential to decrease mask manufacturing costs since it completely eliminates the need for the chrome layer. Instead only π -phase-shifting regions are introduced on a transparent unshifted substrate. Since dark features in CPL can be generated only using destructive interference, the mask pattern may be significantly different from the target pattern. Figure 2.41 demonstrate two examples on how to create a dense line/space pattern using CPL. On the left, the dark features are generated at the edges of the shifted and the non-shifted phase regions. Each phase-shifted region can hence be said to produce two intensity minima, at its phase edges. As a consequence, the width of the line does not determine the width of the dark target. This property allows for imaging of considerably smaller line/space patterns compared to binary masks. Moreover, because of the pronounced intensity drop below the dark lines that is due to the phase interaction, a steeper intensity slope and thus a larger dose latitude can be achieved. In addition a slightly improved depth-of-focus has also been demonstrated (Toh et al., 1991a). In fact, this technique has a similar frequency doubling effect as the alternating phase shifting masks (see Section 2.6.3). It therefore also involves similar design issues, most importantly design issues attributed to phase conflicts.

Wider dark features cannot be produced with the phase-edge technique. For these features, a region between two phase-edges has to be defined that is small enough to completely inhibit the transmission of light. The right part of Figure 2.41 depicts a line/space pattern following this principle. Like the phase-edge technique, the phase-shutter regions have the potential to increase the dose latitude of the lithography process. For the same line/space pattern as before, a chrome-produced dark line on a COG would result a less pronounced intensity minimum. Because spill-over light resulting from higher diffraction orders of the adjacent bright spaces and a constructive phase interaction prevent the intensity to drop to values as low as in the dark-grating CPL case. The latter therefore shows a much improved dose-latitude. The phase-shutter CPL masks resemble attenuated phase-shifting mask (see Section 2.6.3) that have a background transmission of one. In contrast to the phase-edge approach, no design restrictions due to phase conflicts have to be taken into account.

Similarly as with AttPSMs, at first glance the lack of opaque (chromium) regions seems to be prohibitive, for example, for the generation of isolated bright spaces. But this can be made up for by small phase-shifted features that are arranged as small-pitch gratings or checkerboards. There the light is directed into the higher diffraction orders, which are not collected by the projection lens. As a result, only little light passes to the wafer. As an alternative, chromium shields can be added, albeit debilitating the advantage of lower manufacturing demands of CPL masks.

One of the problems of CPL using the phase-shutter technique is an extreme magnification of the dark phase-shifted features, which has to be dealt with by size-biasing these elements (Mack, 1991). However, such a size reduction of the dark features leads to a contrast loss in the aerial image. (Chen et al., 2001) report a much less pronounced magnification under high NAs. Moreover, they propose the application of strong off-axis illumination, with which distinctly larger process windows were achieved. In the same study, the authors demonstrated that CPL is amenable to OPC. In their example, line shortening was reduced using hammerheads. The authors also advance a technique they refer to as the binary half-tone principle.¹⁹ Comparably to the generation of dark areas using checkerboards, small phase-shifted elements are laid out in a pitch below the resolution capabilities of the projector. The elements then stop acting individually but collectively as one phase element, as the authors describe it. The background transmission can be adjusted by modifying the size and the pitch of the elements. This method can, for example, be applied as an alternate linewidth control mechanism Broeke et al. (see also 2002). Recently, a group of authors has proposed the use of pixelated CPL masks as variable high-transmission attenuated PSMs in conjunction with an inverse lithography approach (Borodovsky et al., 2008; Schenker et al., 2008; Singh et al., 2008).²⁰ The background transmission is achieved through a sub-resolution arrangement of phase-elements, thus in a way extending the former authors' binary half-toning technique. Coarsely put, a dense sub-resolution pixel distribution leads to a lower transmission PSM while a sparser set-up increases the background transmission. The authors claim that even under defocus, no image shift compared to a traditional AttPSM occurs. In order to asses the

¹⁹See Section 2.6.1 to compare with the *half-tone biasing* OPC technique.

²⁰ILT-related details of this work will be presented later (see Section 7.1.7).




Figure 2.41: Comparison of chromeless PSM techniques: (a) The wide (300 nm in this example) π -phase-shifted regions in a phase-edge technique allow for non-destructive transmission at their centers. Destructive interference only occurs at the edges of the phase transmission. (b) Phase-shutter regions employ significantly narrower phase-shifted features (75 nm in this example) such that due to the destructive interaction the intensity is completely nulled below them. In order to achieve a dense line/space pattern, a pronounced over-bias of the 0-phase-shifted transparent regions has to be implemented. Both examples were simulated for a 4x system with an NA of 0.8 and a wavelength of 193 nm using a Kirchhoff mask approximation. For the phase-edge mask, the partial coherence factor was set 0.5, whereas the shifter-shutter case uses a factor 0.4. All mask dimensions are in wafer scale.

CPL performance, an extensive numerical study using an electro-magnetic field solver was conducted. One of the results was the determination of a mask resolution limit below which features perform a phase and transmission averaging. Based on this observation, the authors subdivided CPL imaging into three domains, which can be summarized as follows:

1. *Resolved by mask and projector optics*
Standard imaging regime.
2. *Resolved by mask but not by projector optics*
Light is directed mostly in high diffraction orders, very little intensity is transferred into the far-field. Exploitable in form of sub-resolution checkerboard or gratings to implement (even large) opaque areas.
3. *Resolved by neither mask nor by projector optics*
An increase of the pitch frequency to even higher values, beyond the near-field resolution limit, leads to a redirection of light into the DC order. Although not leading to any modulation, these types of features can be used to enhance the image, for example, by its ability to introduce a background intensity.

Another finding was a much less pronounced overbiasing of the dark phase-shifted features in comparison to predictions using thin-mask simulations. Moreover, the severe contrast loss for dense line/space patterns previously suggested by Kirchhoff computations were not confirmed by results obtained with rigorous EMF solvers.

The work also impressively demonstrated why despite a number of appealing properties CPL does not yet enjoy a high popularity: It depends on a tremendous computational and design effort. In addition, novel fabrication and inspection facilities are inevitable (Cheng et al., 2008). In order to move (pixelated) CPL approaches to smaller features sizes, new masks materials will also be required (Schenker et al., 2008).

2.6.4 Pattern Doubling

As NGL approaches such as EUVL or parallel e-beam are far from production-ready, intermediate techniques are required to bridge the gap. Not even the use of aggressive OPC techniques in combination with immersion lithography allows for the realization of the 32 nm HP node and below. As shown before, the minimum k_1 factor for dense lines/spaces is 0.25, which under water immersion translates to a minimum line or space width of about 35 nm. This limit, however, applies only to dense features. Smaller isolated or semi-dense features can be generated.²¹ Thus one solution to overcome the theoretical k_1 limit is to split a dense pattern into several patterns that then exhibit a much lower frequency. The so split patterns can then be individually exposed. Currently, pattern doubling techniques are employed or investigated. Depending on the introduction of NGL, some of the following approaches will have to be extended to triple or quadruple the pitch.

Double exposure

Double exposure can be considered the most straight-forward of all double patterning approaches. The wafer stays in the exposure tool while the two successive exposures are performed. Unfortunately, the resist is said to have a “memory”, that is, all intensity contributions of the first exposure have a correspondence in the concentration of photo-activated components of the resist (Byers et al., 2007; Lee et al., 2008). In other words, the resist shows a linear behavior with respect to the intensity. We can illustrate this by representing the resist response by transfer function (r):

$$r(I_1) + r(I_2) = r(I_1 + I_2),$$

²¹Actually, through an application of processes such as line or space shrinkage, or by over- or under-exposure and/or over- or under-development significantly smaller feature sizes are attainable.

where $I_{\{1,2\}}$ are the intensities obtained with the individual exposure steps. Under such a condition, the double exposure approach does not provide any advantage over a single exposure using the combined pattern. In order to attain a non-linear behavior with respect to the individual exposures, i.e.,

$$r(I_1) + r(I_2) \neq r(I_1 + I_2),$$

alternative resist materials have to be applied. Four common approaches have been proposed:

1. *Two-photon resists*: Two-photon materials incorporate photo-acid generators (PAGs) that require the absorption of two photons for the acid generation (Wu et al., 1992; Yablonovitch and Vrijen, 1999). The response to the exposure intensity is thus proportionate to the squared intensity, yielding

$$r(I_1) + r(I_2) \propto I_1^2 + I_2^2$$

for the double exposure process. Although such materials would indeed enable a double exposure process, their availability for 193-nm lithography has not yet been demonstrated (O'Connor et al., 2008). In addition, high exposure rates would be required, degrading the throughput achievable with this technique.

2. *reversible contrast enhancement layer (RCEL)*: A contrast enhancement layer (CEL) is a bleachable dye applied on top of the photoresist (Oldham, 1987; Kim and Levinson, 2007). The initially opaque layer is bleached through continuous exposure until it finally becomes transparent. Only then the energy is incorporated into the resist to activate the photo-acid generation, thus rendering a non-linear behavior.

RCELs can be reset to the initial opaque state, for example, prior to the second exposure step. The resulting response may be stated as follows (Byers et al., 2007):

$$r(I_1) + r(I_2) \propto I_1 e^{\Delta \cdot I_1} + I_2 e^{\Delta \cdot I_2},$$

where the factor Δ is a product of the bleachable absorbance, the bleaching rate and the thickness of the CEL, and the exposure time. This clearly leads to a non-linear behavior. However, CELs pose several issues including the ability to deposit and remove the material without damaging the underlying resist layer. Moreover, recent studies have shown that for low k_1 factors and high NAs the failure of coupling higher spatial frequencies, which than occur in the RCEL, into the resist, render this approach problematic for double exposure (Shao et al., 2010).

3. *Optical threshold materials*: In optical threshold materials, photochemical occur only under the absorption of a threshold exposure (Byers et al., 2007; Lee et al., 2008; Gu et al., 2009). The non-linear response is given by the following relation:²²

$$r(I_1) + r(I_2) \propto (I_1 \dot{-} I_{th}) + (I_2 \dot{-} I_{th}).$$

In comparative simulation studies (ibid.), it was demonstrated that this approach would exhibit a superior behavior. Although, however, this technique is available in thermoresists, currently no comparable optical materials are present.

Double patterning

As an alternative, especially since double exposure materials are not currently available, double patterning processes are being developed. One of the early proposed techniques is termed *litho-etch-litho-etch (LELE)* (Ebihara et al., 2003; Maenhoudt et al., 2004). As shown in Figure 2.42, LELE constitutes the

²²The dot minus operator ($\dot{-}$) is defined for $x, y \in \mathbb{R}^+$:

$$x \dot{-} y := \begin{cases} x - y & \text{if } x > y \\ 0 & \text{otherwise} \end{cases}$$

most straight-forward approach: The initial stack (a) consists of the photoresist, two different hardmasks and the substrate or any pattern target layer, for example, poly-silicon. Hardmask materials differ depending on the application (i.e., the target layer). Options are poly-silicon, silicon dioxide, metal hardmasks such as Titanium Nitride, or carbon hardmasks. The photoresist is often coated on top of a bottom anti-reflective layer (BARC), which for clarity are not depicted in the sketch. In the first sub-sequence the resist is exposed (b) with the first photomask and developed (c). Then the pattern is transferred into the hardmask by etching (d). The resist strip (e) marks the end of the first step. The second sub-sequence begins with a resist (and BARC) re-coat (f), which is then exposed (g) with the second photomask and developed (h). Both the pattern of the first hardmask and the photoresist pattern are etched into the second hardmask (i). The first hardmask and the resist are stripped (j), finally leaving the second hardmask whose pattern then can be transferred further. Alternatives, using a single hardmask have also been proposed (Dusa et al., 2007). In that case, the second hardmask can be thought of as the final pattern target layer, for example, poly-silicon.

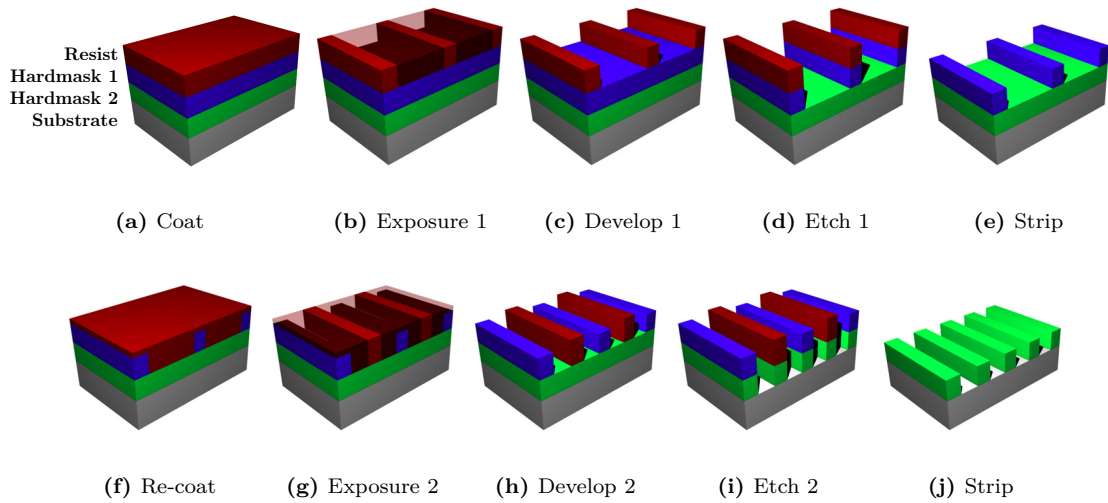


Figure 2.42: Double patterning scheme litho-etch-litho-etch (LELE).

In order to reduce the number of process steps, alternative approaches have been devised, such as *litho-freeze-litho-etch* (LFLE) (Hori et al., 2008) and *litho-litho-etch* (LLE) (Tarutani et al., 2008; Bekiaris et al., 2008; Ando et al., 2008; Nakamura et al., 2010; Maenhoudt et al., 2008). In both techniques, the first etch step is rendered obsolete, in the former case, by chemically protecting the photoresist profile, in the latter one, by stabilizing the resist, for example, by removing excessive solvent from the resist. Both approaches desensitize the first resist such that the profile stays intact after the second develop step. As shown in the litho-freeze-litho-etch schematic in Figure 2.43, after the first sub-sequence (a)–(c), the resist profile is protected (d) so that after the second sequence (e)–(g) both the first and the second profile can be transferred into the hardmask (h).

Self-aligned double patterning

Another approach is called self-aligned double patterning (SADP) (Jung et al., 2006; Bencher et al., 2008)—sometimes also spacer-defined double patterning (SDDP). This approach is well suited for high density lines/spaces arrays, a common situation in memory fabrication. In contrast to the previous techniques including double exposure, it requires only a single lithographic process steps, thus significantly reducing alignment issues. Figure 2.44 shows the schematic process flow. After coat (a), exposure (b) and develop (c), a spacer material is deposited (d).²³ In an successive etch step (e), the excessive

²³Often hardmask materials are used. The deposition can be performed using chemical vapor deposition (CVD).

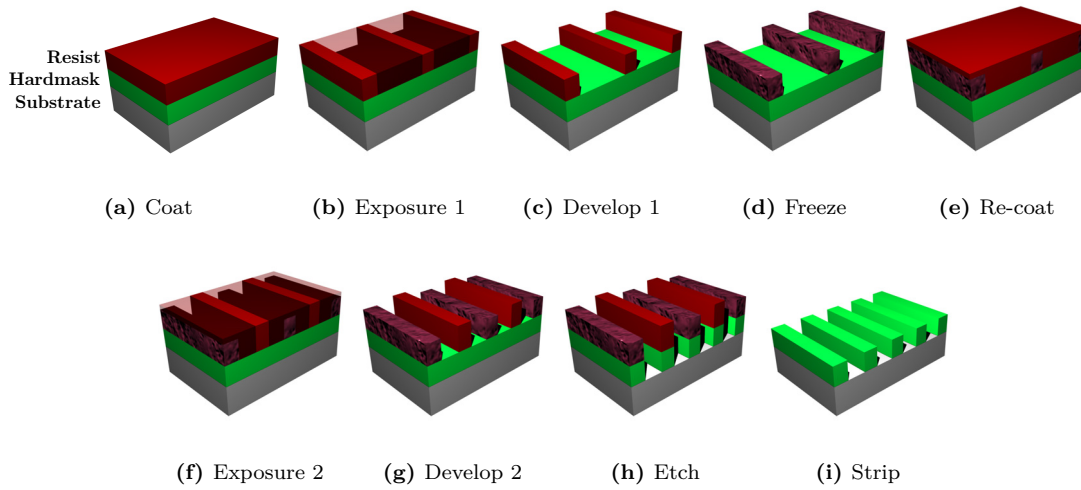


Figure 2.43: Double patterning scheme litho-freeze-litho-etch (LFLE).

spacer material is removed from the top of the resist lines. As a side-effect, the horizontal thickness is also decreased. Thus only the sidewalls of the spacers are left over. Finally, the resist lines are removed (f), leaving only the spacer material, which is then used as an etch mask for the final etch (g). In many applications, a subsequent trimming steps is required, for example, to open up the lines of lines/space pattern.

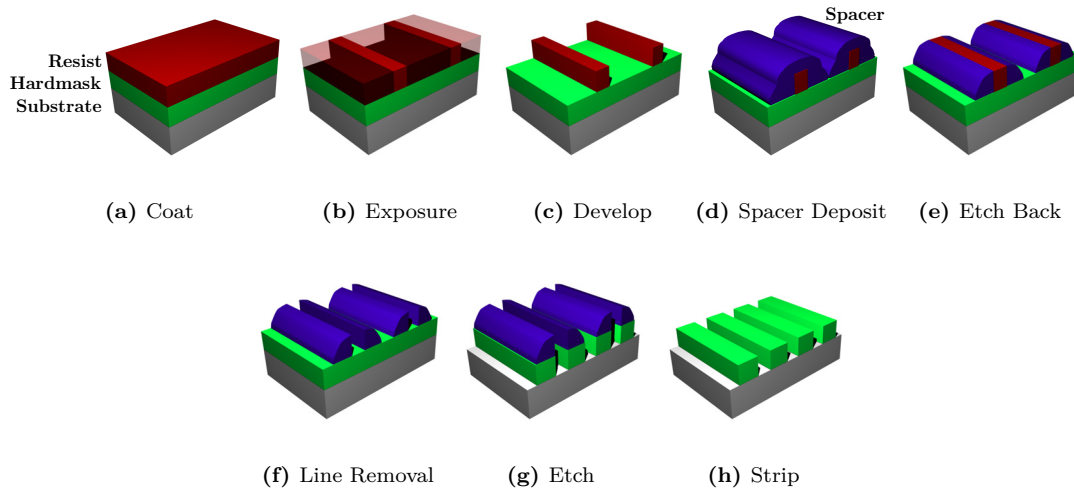


Figure 2.44: Double patterning scheme self-aligned double patterning (SADP).

Dual-tone development

Hinsberg et al. (1989) reported on dual-tone resist (DTR) materials that would incorporate both a positive and a negative respond. Such a material facilitates a self-alignment process. Later, mainly because of abysmal imaging characteristics of this approach, this technique was altered to use a standard

resist material but incorporate a dual-tone development step instead (Asano, 1999). Recent studies demonstrated the general feasibility of the dual-tone development (DTD) approach also for the 193-nm regime, using an aqueous base positive tone and an organic negative tone solvent (Fonseca et al., 2009). Similarly to the spacer-defined double patterning approach, this technique is well suited for lines/spaces layouts. Due to the lower number of process steps it is assumed to exhibit lower costs than SADP.

Issues and challenges

All of the previous options still face numerous technological issues. Therefore, different variants are currently investigated. As many of them exhibit advantages for specific applications and drawbacks for others, it is not unlikely that a combination of options will be applied. Additionally, a combination of approaches, for example, by introducing alternative exposure techniques such as interference lithography have been proposed.

One of the major challenges these techniques entail is posed by their design requirements. Double exposure and double patterning strategies require two photomasks, each of which has to consist of printable features. This has to be considered in the design split. While for regular patterns such as lines/spaces this might be a straight-forward procedure, other components require a more complex treatment. In addition, the design split must not result additional issues, for example, overlay or alignment errors. The sensitivity to alignment errors is reduced with techniques such as SADP or DTD because of the second lithographic process step becomes obsolete. On the other hand, it significantly complicates the IC design. For all approaches, rigorous resolution enhancement techniques are required to reduce overlay and alignment errors and to hence increase the achievable design flexibility.

The cost of ownership associated with any double patterning or double exposure technique is enormous. Research into next generation lithography technologies such as EUVL or parallel e-beam is still conducted with great intense. Once an NGL technology proves its readiness for production, there is a great chance it will replace any operational double patterning approach.

3 Lithography Simulation

Already with its introduction in the 1970s, numerical modeling was an integral part of optical lithography. By the early 1980s, a number of research-oriented and commercial lithography simulators featuring optical and photochemical models were available. Today, lithography simulation is one of the most important tools for lithography process designers and engineers, and is also indispensable in order to incorporate manufacturability considerations into the circuit design. The key role of simulation in lithography is reflected, for example, by the fact that of the technical papers concerned with lithography, only very few do not present simulation results, whereas there is a remarkable number of simulation-only contributions. In all lithographic technology fields, simulation plays an increasingly important role, not only because of an elevated demand due to higher technological complexities but also because the ever-growing computing power provides enhanced simulation potentials.

Lithography simulation can be considered multi-physics in that both the optical and the photochemical aspects of the process have to be taken into account. In this chapter, the optical simulation will be sub-divided into the following three sections:

- (1) The diffraction spectrum resulted from the photomask.
- (2) The transfer function, modeling the behavior of the projection system.
- (3) Wafer-induced effects, including thin-film optics or topography-awareness.

The last section is dedicated to an overview on models for the photochemical processes.

3.1 Photomask Diffraction Spectrum

The diffraction behavior of the photomask can be described using different models with different accuracy domains. A fast approach is to directly employ the Kirchhoff-Fresnel (or Rayleigh-Sommerfeld) diffraction theorem (see Appendix B). This is often termed *thin-mask* or Kirchhoff approach. There, the mask is represented by an infinitesimal thin approximate. Consequently, the diffraction spectrum can be computed by a Fourier transform of the mask function. The photomask can be described, for example, by geometric functions that incorporate phase and transmission values. It is convenient to use primitive geometries such as rectangles or triangles, for which the analytic Fourier transform can be evaluated in a straight-forward manner. As an alternative, a raster-based description of the mask can be used in order to employ a discrete Fourier transform (e.g., the Fast Fourier Transform (FFT) algorithm).

There are a number of limitations that are due to the simplifications assumed when using this Fourier optics approach, most importantly, the underestimation of phase and polarization effects, especially when applied to the computation of phase-shifting masks (e.g., see Erdmann, 2005). In such a regime, the rigorous computation of the electro-magnetic field due to the photomask has to be computed by numerically solving Maxwell's equations (e.g., see Equation (A.1) in Appendix A), for which there are several techniques applicable to computational lithography, including:

- (1) *The Finite-difference time-domain (FDTD)* approach (Yee, 1966),
- (2) *rigorous coupled wave analysis (RCWA)* methods (e.g., Nyysönen and Kirk, 1988; Moharam et al., 1995; Lucas et al., 1996; Turunen, 1996),
- (3) *finite element methods (FEMs)* (Urbach and Bernard, 1989; Bao, 1995; Monk, 2003; Wei et al., 2007; Foulkes and Blaikie, 2008; Burger et al., 2008)
- (4) *finite integration techniques (FITs)* (Weiland, 1977)¹,
- (5) the *boundary element method*, sometimes also called *method of moments* (e.g., Karniadakis, 1989; Monk, 2003), and

¹(also see Rahimi et al., 2010)

(6) the *fast multipole method* (Hafner, 1990; Yeung and Barouch, 1997).

None of the former approaches can be considered absolutely superior. For example, while FEMs are generally well qualified to handle complex geometries, their high memory requirements render it unsuitable for most large scale problems. FDTD approaches, although employing a simplistic geometry representation, are still very popular in lithographic applications because of their moderate memory requirements, the straight-forward fashion in which they allow the extraction of internal fields (a useful property, especially in the scope of rigorous wafer-side computations), and the inherently parallel update scheme. FITs could be considered a compromise between FEMs and FDTD in terms of consumption of resources and the ability to handle complex geometries. However, further thorough studies are required to evaluate their performance in the scope of computational lithography. More recently, RCWA approaches are becoming increasingly popular for computational lithography applications. There the main disadvantage of this approach, the limited ability to represent highly complex geometries, is not significant. Due to the different fields of applications, there are only a few comparative studies (Berger et al., 2004; Erdmann et al., 2006b; Burger et al., 2008; Rahimi et al., 2010), neither of which could clearly expose one approach over another.

In order to employ computational electro-magnetics approaches to large scale problems such as the field computation of large mask areas, the performance of either technique has to be drastically improved. In addition to domain decomposition techniques and parallelization techniques, hybrid approaches as, for example, proposed by Hugonin et al. (2008), may be an answer.

The two most frequently employed EMF solvers for the computation of the field of the photomask, FDTD and RCWA, are briefly discussed in the following paragraphs.

3.1.1 Finite-difference Time-Domain

The FDTD method was first proposed by Yee (1966) and significantly extended by Taflov and Brodwin (1975). It is one of the first solvers that were applied to lithography-specific EMF problems such as computations of the electro-magnetic field of the photomask or of the wafer topography (Wojcik et al., 1987; Guerrieri et al., 1991; Wong and Neureuther, 1995; Erdmann and Friedrich, 2000). The FDTD method makes use of a decomposition of Maxwell's curl equations [(A.1a) and (A.1b)] into the individual electric and magnetic field components. The simulation domain is represented by a equidistant Cartesian grid, where the electric and magnetic field components are staggered on the vertices, shown in Figure 3.1.




Figure 3.1: FDTD unit cell: electric and magnetic field components are staggered on the vertices (Yee, 1966).

In addition to the discrete spatial domain, discrete time steps are used. All functions of space and time are evaluated on the discrete spatial and temporal domain rather than on the continuous coordinates. In order to guarantee for a numerically stable exploitation of the dependence of the electric field components on the magnetic field components, a staggered update scheme is employed. In other words, in half a time step the electric field components are updated, whereas the updated magnetic field is computed after a full time step (or vice versa). For example, given a time step of duration Δt and

a cell size of Δx ,² and cell coordinates, denoted by $[i, j, k]^T$, the x component of the magnetic field (H_x) is assumed to be constant in the spatial range $[i - 1/2, i + 1/2]$ and within the temporal range $[t - 1/2, t + 1/2]$. Consequently, the electric field component (E_x) shows a constant behavior for $[i, i + 1]$ and $[t, t + 1]$.

As common to finite-difference approaches, the selection of the boundary conditions of the simulation area is a critical determinant. For example, to emulate an open region, absorbing layers have to be introduced. Various related techniques have been proposed (for a comparative study, see Blaschak and Kriegsmann, 1988), none of which, however, are fully suited to suppress reflections under oblique incidence angles. State-of-the-art FDTD implementations often employ a technique proposed by Berenger (1994), and which he termed perfectly matched layer (PML). In contrast to former matched-layer-approaches, both the electric and magnetic field intensities, and the electric and magnetic conductivities are sub-divided with respect to the incidence direction, consequently leading to an extension of Maxwell's curl equations. The additional degree of freedom allows the layers to be adjusted to behave neutrally in virtually all circumstances and hence to significantly reduce back-reflections.

The discretization both of the spatial and the temporal domains have to be selected according to the problem: Typically, the cell sizes have to be in the order of a fraction of the wavelength. Similarly, the time step size has to be chosen in accordance with the wavelength of the light and the spatial frequency of the structures. Another issue, especially in lithography, where the steady state rather than initial or intermediate states is of interest, is the determination of the number of time steps. Often, it can be determined by estimating the number of time steps required for a wave to be propagated through highly reflective or refractive regions.

Numerous extensions to the standard FDTD approach have been proposed, many of which are concerned with exploiting its inherent parallelism (e.g., Guerrieri et al., 1991; Adams et al., 2007)—all cells can be updated concurrently. In addition, other general finite-difference techniques such as the alternating direction implicit method have been successfully applied to FDTD (Namiki, 1999). Computation times and memory requirements inhibit the three-dimensional computation of large mask or wafer areas, even using state-of-the-art parallelization techniques. Adam (2001) has hence proposed a domain decomposition technique in which a three-dimensional problem is divided into multiple two-dimensional sub-problems. The so obtained fields are then superposed,³ exploiting the linearity of the Kirchhoff-Fresnel formulation. With this approach, large-scale problems that were previously not computable can be solved. The obtained results are approximate solutions. The validity must therefore be considered on a case-by-case basis.

3.1.2 Rigorous Coupled Wave Analysis and The Waveguide Method

As an alternative EMF solver, the RCWA method, which is a Fourier modal technique, gains increasing attention in the scope of computational lithography (Nyyssonen and Kirk, 1988; Moharam et al., 1995; Turunen, 1996; Lucas et al., 1996). In this work too, EMF simulations are conducted using an RCWA approach, called the *waveguide method* (Nikolaev and Erdmann, 2003; Evanschitzky and Erdmann, 2005). In contrast to FDTD, both the electro-magnetic field and the permittivity profile are represented by Fourier series expansions.

An arbitrary 3-D photomask geometry (schematically shown in Figure 3.23.2(a)), which may consist of multiple absorbers and the mask substrate (or a wafer topography definition), is subdivided into layers that are homogeneous in the optical axis direction (z -direction; Figure 3.23.2(b)). Because of the z homogeneity, each layer can then be described by a two-dimensional geometry and the respective material properties (Figure 3.23.2(c)). The elimination of the z dependence of the material properties plays a central role in RCWA approaches as it allows for the representation of the materials in only lateral dependence as a Fourier series. This property is sometimes referred to as the *waveguide assumption*. The selection of an appropriate layer decomposition strongly impacts the accuracy of the result. For a feasible numerical treatment, the layers are then decomposed into two-dimensional geometric

²In general, the cells do not need to be equally sized in all directions. For practical numerical reasons, however, this is often the case.

³subtracting the one-dimensional intersection to avoid their redundant contribution




Figure 3.2: Photomask representation of the waveguide method: (a) initial 3-D topography, (b) decomposition into z -homogeneous slices, (c) 2-D representation of one slice, (d) slice after rectangulation.

primitives, typically rectangles, which also allow for an optimized treatment of the TM polarization mode as proposed by Li (1996); Lalanne and Morris (1996); Granet and Guizal (1996). An example rectangulation is shown in Figure 3.23.2(d).

Both the material property distribution and the electric and magnetic fields are represented by Fourier series expansions up to a user-defined order M . For each layer, an eigenvalue problem for the material and field series is laid down, solving for both the propagating and evanescent modes inside the layer. The layers are then coupled utilizing appropriate boundary conditions (e.g., a scattering matrix approach), finally resulting in the plane wave distribution on top and at the bottom of the simulation domain (e.g., the photomask area).

Commonly, periodic boundary conditions are applied in the x/y direction of each layer. Thus similarly to other approaches including the thin mask model, isolated features have to be emulated by defining the period such that a negligible cross-talk between periodic features occurs.

In addition to the top and bottom spectra, the spectra for the individual layers can be extracted. This is required, for example, to compute the light distribution inside inhomogeneous resist (and anti-reflective coating) layers; a task typical for pattern doubling applications.

The time complexity of RCWA methods is

$$O(n \cdot m^3),$$

where n is number of layers into which the problem is sliced and m is the product of the employed orders of the Fourier series in x and y directions. Like the number of layers, the required number of orders depends on the problem. For example, Evanschitzky et al. (2007) proposes the following estimation rule for optical

$$M_x \approx \frac{6p_x}{\lambda} + 1 \quad M_y \approx \frac{6p_y}{\lambda} + 1,$$

and similarly for EUV photomasks

$$M_x \approx \frac{p_x}{\lambda} + 1 \quad M_y \approx \frac{p_y}{\lambda} + 1,$$

where p_x and p_y denote the respective x and y period.

In a similar fashion as proposed by Adam (2001) for the FDTD method, a domain decomposition for the waveguide method was employed by Evanschitzky et al. (2007) and later extended to a concurrent version by Shao et al. (2008). The three-dimensional domain is decomposed in the following manner:

- (1) The simulation area is partitioned with cuts parallel to the x axis. Each partition is to be homogeneous (with respect to the dielectric material properties) in y . A two-dimensional x/z slice that represents the so obtained partition is stored.

(2) In the same way, y/z cuts are generated.

(3) In addition, the one-dimensional intersections between the x/z and the y/z cuts are stored.

For each of the so obtained slices, a 2-D or 1-D computation, respectively, is conducted. The resulted spectra are then used to reconstruct the 3-D spectrum. This is done in the frequency domain by expanding the 2-D cuts through appropriate filters and a subsequent superposition of the individual parts. Prior to that, in order to prevent double accounting for the intersections between the x/z and the y/z cuts, the 1-D cuts have to be subtracted accordingly. The drastic performance increase both in terms of time complexity and memory consumption compared to fully three-dimensional computations has been amply demonstrated, and in many applications, the results closely approximate those obtained with a full 3-D approach (Evanschitzky et al., 2007). However, exact validity constraints are yet subject to further research.

3.2 Imaging

In order to predict the image obtained by the lithographic projection system, its transfer function has to be determined. This can be done in a rigorous manner, for example, by optical design models. In most applications of lithography simulation, however, the projection lens system is regarded as a black box, shown in Figure 3.3.




Figure 3.3: Imaging system.

As presented in Appendix C, the scalar imaging equation under partially coherent illumination can be stated as follows:

$$\begin{aligned}
 I(x_1, y_1) &= J_I((x_1, y_1), (x_1, y_1)) \\
 &= \iiint\limits_{-\infty}^{\infty} J_C(x_0 - x'_0, y_0 - y'_0) O(x_0, y_0) O^*(x'_0, y'_0) \\
 &\quad H(x_1 - x_0, y_1 - y_0) H^*(x_1 - x'_0, y_1 - y'_0) dx_0 dy_0 dx'_0 dy'_0,
 \end{aligned} \tag{3.1}$$

where O is the object function—the field⁴ of the photomask in the lithography case— H is the projector transfer function, and J_C is the mutual intensity, a weight factor, of two points under extended source

⁴either assuming a Kirchhoff approximate or obtained through rigorous EMF calculations

conditions such as Köhler or critical illumination. A brief interpretation Appendix C can be stated as follows: The intensity at a point in the image plane is given by the propagation of the mutual intensity of all contributing points, that is, of all points that lay in the support of the projection system and the illuminator.

As we shall see later, it is often convenient to represent the former equation in the spatial frequency domain. This allows for a decomposition into a sum of incoherent source points (see Section 3.2.5), or for an encapsulation of the illuminator and projector contribution and its separation from the impact of the mask (see Section 3.2.6). By using the Fourier integral representation of individual components J_C , H and O , we have:

$$\begin{aligned} I(x_1, y_1) = & \int \int \int \int \int \int_{-\infty}^{\infty} \mathcal{F}(J_C)(f, g) \mathcal{F}(O)(f', g') \mathcal{F}(O)^*(f'', g'') \\ & \mathcal{F}(H)(f + f', g + g') \mathcal{F}(H)^*(f + f'', g + g'') \\ & \exp(-i2\pi((f' - f'')x_1 + (g' - g'')y_1)) df dg df' dg' df'' dg''. \end{aligned} \quad (3.2)$$

3.2.1 Aberrations

An ideal diffraction-limited lithographic projection system would yield a plane wavefront in the focal plane. As already argued in Section 2.2, however, there are no aberration-free projection systems. The deviation of the actual wavefront from the ideal one can be defined as a function Φ , which can be expressed as a series of polynomials (Zernike, 1934):

$$\Phi(\rho, \theta) = \sum_i c_i Z_i,$$

where c_i are the coefficients of the individual polynomials Z_i , called Zernike polynomials. The parameters ρ and θ are the normalized polar coordinates in the pupil domain. The Zernike polynomials can be considered a derivation of Jacobi's polynomials in the circular domain and are hence orthogonal over the unit circle (the pupil). They have the following form:

$$Z_j(\rho, \theta) := \begin{cases} R_n^m(\rho) \cos(m\theta) & \text{if } j \text{ is even or } m = 0 \\ R_n^m(\rho) \sin(m\theta) & \text{if } j \text{ is odd} \end{cases}$$

The radial polynomials R are defined as:

$$R_n^m(\rho) := \sum_{s=0}^{\frac{n-m}{2}} (-1)^s \frac{(n-s)!}{s! \left(\frac{n+m}{2} - s\right)! \left(\frac{n-m}{2} - s\right)!} \rho^{n-2s}.$$

For the integers m and n , the following statements are frequently used (for example, in optics design software; see (Wong, 2005, chap. 7, p. 137))

$$\begin{aligned} & n \geq m \\ & n - m \text{ is even} \\ & n := \left\lceil \frac{1}{2} \left(-3 + \sqrt{1 + 8j} \right) \right\rceil \\ & m := n - 2 \left\lfloor \frac{1}{4} ((n+2)(n+1) - 2i) \right\rfloor \end{aligned}$$

Depending on the application, different normalizations such as root-mean-squares can be additionally introduced.

3.2.2 Vector Imaging

In contrast to the previous assumption, light is a vectorial rather than a scalar entity. Especially for high numerical apertures, polarization effects have to be considered, and consequently the vector nature of light has to be taken into account. Since the photochemical reaction are due to the electric field,⁵ and because magnetic and electric fields are perpendicular to each other, it is convenient to restrict the formal analysis to the electric field. Furthermore, the extension of the scalar diffraction theory to a vectorial one is straightforward in that the respective equations apply to all components of the electric field.

Within the optical system, the propagation direction is along the optical axis, which we define as z . For the computation of the pupil function, the z component can hence be omitted.⁶ Moreover, in the pupil domain the x and y components are identical to the s and p polarization components. Outside the pupil, the propagation direction is generally different from z such that the field components have to undergo a transformation according to the new direction. We will refer to this transformation as polarization correction. Within the lithographic projection system, two polarization corrections are required (depicted in Figure 3.4):

- (1) The incident plane waves have to be transformed to the normal direction in the pupil.
- (2) At the exit pupil, the field components have to be reconstructed according to the newly arising propagation direction.




Figure 3.4: Polarization coordinates transform: Two polarization corrections are required when evaluating the wave propagation through the optical system: (1) when the light enters the pupil function (propagation direction vector $\hat{\mathbf{k}}_1$ and (2) when the waves emerge from the exit pupil toward the output plane ($\hat{\mathbf{k}}_3$). Within the pupil, light propagation is assumed to be in the z direction ($\hat{\mathbf{k}}_2$). Therefore the x component is equal to the p polarization mode, and the y component coincides with the s mode.

The extension of the scalar imaging theory to a vectorial one has been subject of numerous researches (for example, Yeung, 1988; Flagello and Milster, 1992; Gallatin, 2001; Flagello et al., 2005; Evanschitzky et al., 2009). For the description of the polarization correction we will use a formalism first proposed by Mansuripur (1986). Let us first consider the exit pupil case (2) in which a TE/TM field in the optical system is to be transformed into x , y and z components in the output plane. Given an electric field in the optical system

$$\tilde{\mathbf{E}} = \begin{pmatrix} E_x \\ E_y \end{pmatrix} = \begin{pmatrix} E_p \\ E_s \end{pmatrix}$$

⁵Experimentally, this was first demonstrated by Wiener (1890)

⁶For a given propagation direction, the z component is actually determined implicitly because of the orthogonality of the electro-magnetic field components.

and the unit propagation vector

$$\hat{\mathbf{k}} = \begin{pmatrix} k_x \\ k_y \\ k_z \end{pmatrix} \in \mathbb{R}^3.$$

The target is then to find a mapping $\mathbf{M} : \mathbb{C}^2 \rightarrow \mathbb{C}^3$ that transforms the field components according to the new propagation direction,

$$\mathbf{E} = \mathbf{M}\tilde{\mathbf{E}}.$$

In fact the following matrix $\mathbf{M} \in \mathbb{R}^{3 \times 2}$ provides the required property:

$$\mathbf{M} = \begin{bmatrix} m_{xx} & m_{yx} \\ m_{xy} & m_{yy} \\ m_{xz} & m_{yz} \end{bmatrix} := \begin{bmatrix} 1 - \frac{k_x^2}{1 + k_z} & -\frac{k_x k_y}{1 + k_z} \\ -\frac{k_x k_y}{1 + k_z} & 1 - \frac{k_y^2}{1 + k_z} \\ -k_x & -k_y \end{bmatrix},$$

where k_x, k_y, k_z are the components of $\hat{\mathbf{k}}$ (i.e., $\hat{\mathbf{k}}_3$). We shall refer to this matrix as Mansuripur matrix.

The first case (1), the transition from the incidence direction to the normal direction in the optical system, can be restricted to a $\mathbb{C}^2 \rightarrow \mathbb{C}^2$ mapping since the mask spectrum is fully determined by its x, y components. The transformation can hence be conducted using the inverse of the partial Mansuripur matrix whose last row (the z transformation) is omitted:

$$\mathbf{M}' := ((m_{ij}))_{2 \times 2}.$$

The inverse of \mathbf{M}' does not exist if and only if the considered wave is evanescent.⁷ That is, the propagation vector has to have a non-zero z entry:

$$|\mathbf{M}'| \neq 0 \iff k_z \neq 0 \quad (\text{note, } \hat{\mathbf{k}} \in \mathbb{R}^3).$$

Because of the unity of $\hat{\mathbf{k}}$, the k_z component is given by:

$$k_z = \sqrt{1 - (k_x^2 + k_y^2)}.$$

With the transformation matrix laid out, we can use the following relation for the coordinate system transformation:

$$\tilde{\mathbf{E}} = \mathbf{M}'^{-1} \mathbf{E}', \tag{3.3}$$

where \mathbf{E}' is the incident field and $\tilde{\mathbf{E}}$ is again the field in the pupil.

3.2.3 Jones Calculus

In addition, aberrations, especially polarization effects of the projection system, are commonly described using a calculus proposed by Jones (1941).⁸ As before, the light propagation is assumed to be along the z -axis. The amplitude and phase of the electric field can then be expressed as a vector \mathbf{E} where the components constitute the x and y contributions—or, alternatively, the TE and TM modes. Polarization actions can then be modeled as a matrix \mathbf{M} describing linear transformations:

$$\begin{aligned} \mathbf{M} : \mathbf{E} &\mapsto \mathbf{E}' \\ \mathbf{E}' &= \mathbf{M}\mathbf{E} \end{aligned}$$

⁷Non-propagating waves are obviously not processed by the imaging model.

⁸An exhaustive discussion on the application of Jones' calculus and Stokes' parameters (see next section) to the lithography case was led by Flagello et al. (2005).

with

$$\mathbf{M} := \begin{bmatrix} m_{xx} & m_{xy} \\ m_{yx} & m_{yy} \end{bmatrix}$$

For example, given the electric field along the propagation direction as:

$$\mathbf{E} := \begin{pmatrix} E_x \\ E_y \end{pmatrix},$$

a linear horizontal polarizer (resulting a purely p -polarized state) can be expressed in terms of \mathbf{M}' , with

$$\mathbf{M}' = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}.$$

3.2.4 Stokes Parameters

An alternative representation has been developed by Stokes (1852). His parameters can be considered an implicit formulation of polarizing elements. In contrast to the Jones formalism, in which completely polarized light is assumed, the Stokes parameters stochastically capture electric field variations of quasi-monochromatic and partially polarized or unpolarized light. This formulation is hence well suited to incorporate measured data, for example, of the illumination source, into simulation. The Stokes parameters can be defined in the following vector form:

$$\begin{pmatrix} s_0 \\ s_1 \\ s_2 \\ s_3 \end{pmatrix} = \begin{pmatrix} \langle a_x \rangle^2 + \langle a_y \rangle^2 \\ \langle a_x \rangle^2 - \langle a_y \rangle^2 \\ 2\langle a_x a_y \cos(\delta) \rangle \\ 2\langle a_x a_y \sin(\delta) \rangle \end{pmatrix},$$

where a_x and a_y are the amplitudes of the x and y components of the electric field, and δ is the phase difference. In the lithography case, the time average is often omitted. That is, the waves are approximated as being time-harmonic. The first component expresses the total intensity, the last three components can be thought of as the intensity response to the following polarization states:

- s_1 (also Q) describes the preponderance of vibrations in x direction over ones in the y axis.
- s_2 (also U) describes the same preference as s_1 , however, in a $\pi/4$ -rotated coordinate system.
- s_3 (also V) constitutes similar preference response as s_1 and s_2 in a polar coordinate system. It hence expresses the preponderance of the polarization to be either clockwise or counterclockwise.

An equivalent polarization description can be derived from the polarization ellipse, shown in Figure 3.5. At a given locus, the electric field components can be illustrated as a function of time. For any polarization state, the graph describes an ellipse. Here a line—resulted in the case of linear polarization—is regarded as a deformed ellipse. The polarization ellipse is commonly specified by the azimuth angle of the ellipse (the polarization angle ψ), and the amplitudes of the wave components (a_1 and a_2). Alternatively, the parameters of the ellipse can be given in spherical coordinates, as points on a Poincaré sphere (see Figure 3.6). There, in addition to the polarization angle, the ellipse is defined by its ellipticity (χ) and the radius of the sphere, which corresponds to the degree of polarization (p), i.e., the ratio of polarized as opposed to unpolarized light. A polarization state given in Stokes' parameters can be converted into the Poincaré notation as follows:

$$I = s_0 \tag{3.4a}$$

$$p = \frac{\sqrt{s_1^2 + s_2^2 + s_3^2}}{I} \tag{3.4b}$$

$$2\psi = \arctan \frac{s_2}{s_1} \tag{3.4c}$$

$$2\chi = \arctan \frac{s_3}{\sqrt{s_1^2 + s_2^2}}. \tag{3.4d}$$




Figure 3.5: Polarization ellipse: Any polarization state can be determined by an ellipse that is specified by the bounding box resulted from the amplitudes of the wave components (a_1 and a_2) and its azimuth or tilt angle (ψ). Alternatively, the ellipticity can be specified as the angle between the half-height and the half-width of the ellipse, conveniently determined in the tilted coordinate system of the ellipse (ξ, η).




Figure 3.6: Pointcaré sphere: each point on the sphere corresponds to polarization state P . The radius of the sphere denotes the degree of polarization, that is, the ratio between polarized and unpolarized light. The points s_1 , s_2 and s_3 define Stokes' parameters.

3.2.5 Abbe's Approach

The previous considerations lay the foundation for the simulation of *aerial images*,⁹ that is, the intensity distribution resulted at the pupil plane. In order to formulate an according algorithm, an adequate numerical representation of the former equations is required. In this section, the Abbe approach, which is the image formation method used in this work, and its extension to a fully vectorial imaging approach are presented. The following section introduces an alternative method—often called Hopkins' approach and traditionally used in OPC applications—and a comparison of both methods.

The Abbe approach makes use of the fact that illumination sources used in projection systems such as microscopy (and also lithography) can be said to consist of mutually incoherent source points (Abbe, 1873, Born and Wolf, 1999, chap. 8, p. 467). Mathematically this can be expressed as integration over the source (Wong, 2005, chap. 8, pp. 155–157). The arguments of the source integral are hence the coherent terms resulted from one source point and the field of the mask, imaged by the transmission function. In the scalar imaging case, this can be derived as follows.

Given the scalar imaging equation (3.2), the mutual intensity (J_C) solely depends on the frequency f, g and can hence be factored out, yielding,

$$I(x_1, y_1) = \iint_{-\infty}^{\infty} \mathcal{F}(J_C)(f, g) \left(\iiint_{-\infty}^{\infty} \mathcal{F}(O)(f', g') \mathcal{F}(O)^*(f'', g'') \right. \\ \left. \mathcal{F}(H)(f + f', g + g') \mathcal{F}(H)^*(f + f'', g + g'') \right. \\ \left. \exp(-i2\pi((f' - f'')x_1 + (g' - g'')y_1)) \, df' dg' df'' dg'' \right) df dg. \quad (3.5)$$

The fourfold integral of (3.5) can now be separated into a double integral that depends only on $(f, g; f', g')$ and one dependent only on $(f, g; f'', g'')$,

$$I(x_1, y_1) = \iint_{-\infty}^{\infty} \mathcal{F}(J_C)(f, g) \left(\iint_{-\infty}^{\infty} \mathcal{F}(H)(f + f', g + g') \mathcal{F}(O)(f', g') \right. \\ \left. \exp(-i2\pi(f'x_1 + g'y_1)) \, df' dg' \right) \\ \left(\iint_{-\infty}^{\infty} \mathcal{F}(H)^*(f + f'', g + g'') \mathcal{F}(O)^*(f'', g'') \right. \\ \left. \exp(i2\pi(f''x_1 + g''y_1)) \, df'' dg'' \right) df dg. \quad (3.6)$$

By setting

$$A'_{(f,g)}(f', g') := \mathcal{F}(H)(f + f', g + g') \mathcal{F}(O)(f', g')$$

and noting that the inner double integrals of (3.6) then constitute the Fourier transform (and the complex conjugate Fourier transform, respectively) of A' , we obtain,¹⁰

$$I(x_1, y_1) = \iint_{-\infty}^{\infty} \mathcal{F}(J_C)(f, g) \mathcal{F}(A'_{(f,g)})(x_1, y_1) \mathcal{F}(A'_{(f,g)})^*(x_1, y_1) \, df dg \\ = \iint_{-\infty}^{\infty} \mathcal{F}(J_C)(f, g) \left| \mathcal{F}(A'_{(f,g)})(x_1, y_1) \right|^2 \, df dg. \quad (3.7)$$

Numerically (3.7) has to be solved for a discrete number of mutually independent source points. Regardless of the discretization technique, the source can hence be represented as a set of source points

⁹They are sometimes also called coupling images, or, in the immersion case, aquatic images.

¹⁰Often the inverse Fourier transform is used in order to avoid mirroring of the coordinate system. Also from the systems point of view an inverse transform seems more appropriate since the result is in the spatial domain.

$\{(f_\sigma, g_\sigma; j_\sigma)\}$, where each source point defines a pair of spatial frequencies and its discrete magnitude. The total intensity in the image plane can then be written as,

$$I(x_1, y_1) = \sum_{\sigma} j_{\sigma} |a_{\sigma}(x_1, y_1)|^2, \quad (3.8)$$

where the amplitude of the coherent term, i.e., the object spectrum and the transmission function, is given by:

$$a_{\sigma}(x_1, y_1) := \mathcal{F}(A'_{(f_{\sigma}, g_{\sigma})})(x_1, y_1).$$

For clarity's sake, the Abbe approach—sometimes also referred to as source points integration method—is sketched here for the scalar imaging case. As shown earlier (see Section 3.2.2), it can be extended to vector imaging in a straightforward way. The frequency dependence of the amplitude of the coherent term then serves another purpose, additionally to the offset to be employed for the transmission function: it allows for the selection of an incidence-angle-specific object spectrum that has been evaluated, for example, through rigorous EMF simulation. This option provided by the Abbe approach constitutes a significant difference compared with the techniques presented later on. On the other hand, since the optical system consisting of the illuminator and the transmission function has to be re-evaluated even if only the object has changed, the computation time may, under certain conditions, be magnitudes larger than for alternative techniques. Particularly the computation time linearly depends on the number of source points used for the computation. The selection of the source sampling technique is thus a critical property that controls the trade-off between computation time and accuracy.

The impact of the polarization state of the illumination source can be decomposed into the s (TE) and p (TM) contributions. For that purpose all steps in the computation of the farfields are conducted separately for the individual polarization modes. In each step, the system is assumed to be fully TE or TM polarized, respectively. Of course, that also requires the input spectrum (or electro-magnetic field) to be given in a TE/TM component representation. The intensity in the image plane arising from the actual polarization setting of a single source point is then given as follows: While the individual $(x, y$ and $z)$ field components are incoherent, the amplitudes of the polarization components have to be superposed coherently. Moreover, the degree of polarization leads to two terms that have to be considered independently and combined into a weighted sum. Formally,

$$I(x, y) := (1 - d) \left(\sum_{\mathbf{c} \in \{x, y, z\}} \frac{1}{2} |E_{\mathbf{c}}^s(x, y) e^{i\delta}|^2 + \frac{1}{2} |E_{\mathbf{c}}^p(x, y)|^2 \right) + d \left(\sum_{\mathbf{c} \in \{x, y, z\}} |E_{\mathbf{c}}^s(x, y) e^{i\delta} \sin \psi + E_{\mathbf{c}}^p(x, y) \cos \psi|^2 \right), \quad (3.9)$$

where d is the degree of polarization, δ is the polarization phase difference and ψ is the polarization angle for a given source point.¹¹ The superscripts s and p denote the respective polarization modes.

¹¹This representation is equivalent to Stokes' parameters (see Section 3.2.4).

Equations (3.8) and (3.9), and the vector imaging considerations in Section 3.2.2 lead to the following Abbe imaging algorithm (Evanschitzky et al., 2009).

Algorithm 3.1: Schematic Abbe vector imaging approach.

```

Data:  $\Sigma$ ; // set of source points
Data:  $\{\mathcal{F}(\mathbf{O})_\sigma\}$ ; // mask spectra, (single spectrum for Hopkins assumption)
1 foreach source point  $\sigma \in \Sigma$  do
2    $\mathcal{F}(\mathbf{O})_\sigma \leftarrow$  retrieve mask spectrum;
3   foreach  $\mathbf{E}[f_i, g_i] \in \mathcal{F}(\mathbf{O})_\sigma$  do // e-field with  $x, y$  components for each polarization
4     mode:  $\mathbf{E}[:, :] = [\mathbf{E}^s[:, :], \mathbf{E}^p[:, :]] \in \mathbb{C}^{2 \times 2}$ 
5     if frequency  $f_i, g_i$  not collected by projector then
6       continue;
7     for  $m \in \{s, p\}$  do // for both polarization modes
8        $\mathbf{E}^m[f_i, g_i] \leftarrow \mathbf{M}'^{-1} \mathbf{E}^m[f_i, g_i];$  //  $\mathbf{M}' \in \mathbb{R}^{2 \times 2}$  (mansuripur)
9       if focus  $\neq 0$  selected then
10         $\mathbf{E}^m[f_i, g_i] \leftarrow \mathbf{E}^m[f_i, g_i] \phi;$ 
11        // vectorial pupil aberrations (Jones matrix:  $\mathbf{J}[:, :] \in \mathbb{R}^{2 \times 2}$ )
12         $\mathbf{E}^m[f_i, g_i] \leftarrow \mathbf{J}[f_i, g_i] \mathbf{E}^m[f_i, g_i];$ 
13        // scalar pupil aberrations (e.g., Zernike  $\Phi[:, :] \in \mathbb{C}$ )
14         $\mathbf{E}^m[f_i, g_i] \leftarrow \mathbf{E}^m[f_i, g_i] \Phi[f_i, g_i];$ 
15        // image plane spectrum ( $x, y, z$  components):  $\mathbf{E}'[:, :] = [\mathbf{E}'^s[:, :], \mathbf{E}'^p[:, :]] \in \mathbb{C}^{3 \times 2}$ 
16         $\mathbf{E}'^m[f_i, g_i] \leftarrow \mathbf{M} \mathbf{E}^m[f_i, g_i];$  //  $\mathbf{M} \in \mathbb{R}^{3 \times 2}$  (mansuripur)
17      // retrieve fields by performing Fast Fourier Transform (FFT) for each
18      // component for each polarization mode
19      for  $c \in \{x, y, z\}$  do
20        for  $m \in \{s, p\}$  do
21           $\hat{E}_c^m \leftarrow \text{FFT}^{-1}(E_c^m);$ 
22      // superposition of individual polarization contributions and fields
23      // (amplitudes) [see (3.9)]
24       $I_\sigma \leftarrow \bigoplus_{c \in \{x, y, z\}, m \in \{s, p\}} \hat{E}_c^m$ 
25      // incoherent sum over source points yields final image
26  $I \leftarrow \sum_\sigma I'_\sigma$ 

```

The first part of the input data in Algorithm 3.1 is the set of source points, obtained through an arbitrary sampling technique (e.g., a regular source grid). As before, each source point is defined by its spatial frequency, the intensity and phase, and the polarization state, for example, using Stokes' convention (see Section 3.2.4). The second input data set consists of the discrete photomask field in the spatial frequency domain. If the spectrum is assumed to be independent¹² of the incidence angle of the light (Hopkins' assumption), a single spectrum, evaluated at the normal incidence, is provided. The Kirchhoff (thin-mask) approach implies this assumption. If spectra for multiple incidences are available—typically through rigorous EMF computations (see Section 3.1)—they are provided, for example, in form of a look-up table. Spectra for missing inclination angles can be complemented by inter- or extrapolation. In Algorithm 3.1, the spectrum corresponding to a source point σ is denoted by $\mathcal{F}(\mathbf{O})_\sigma$.

The algorithm can be briefly characterized as follows:

- (1) The outer loop (l. 1) iterates over the source points. The result of each step is a source-point-induced intensity field. The sum of all fields constitutes the total intensity distribution.
- (2) Within a second inner loop (l. 3), the spatial object frequencies are propagated through the projection system. The s and p polarization modes are treated independently (here indicated as another

¹²aside from a shift

loop [l. 6]). As a first step, the field components are converted from the angle of incidence to the normal incidence direction within the projection system, using the Mansuripur formalism. The actual pupil functions including de-focus, and scalar and vectorial aberrations are then applied. The field components are updated according to the propagation direction resulted from the light emerging from the exit pupil.

- (3) The so obtained spectra are back-transformed into the spatial domain (l. 13), and the resulted fields are superposed [see (3.9)] to yield the source-point-induced intensity distribution (l. 15).

The Abbe approach does not depend on a specific representation (e.g., discretization) of the source, the pupil or the image plane. Common implementations employ rectangular grids on all three components. This is especially convenient if user data such as measured sources or pupils are considered. In addition it allows the accuracy to be adjusted a priori. Recently, we have proposed an alternative sampling technique, allowing for arbitrarily positioned source points and using a continuous projector coordinate system (Evanschitzky et al., 2009). We have demonstrated that with this technique, even with a relatively sparse sampling density a high accuracy can be maintained.

Another variant in which the pupil function is decomposed into kernels—each of which corresponds to one source point—was recently presented (Yamazoe, 2008; Chen et al., 2008; Lian and Zhou, 2009).¹³ Because the employed eigen-decomposition (using a singular value decomposition (SVD)) is performed on a matrix, this approach depends on a rectangular grid representation of the pupil. Moreover, the same number of Fourier transforms, which account for the major fraction of the total computation time, as in a conventional implementation is required. The approach still presents a number of advantages over conventional Abbe or Hopkins approaches:

- Once precomputed, the eigenfunctions present a numerically efficient alternative to the conventional computation of the pupil function: It is reduced to a plain multiplication of the object spectrum and the kernels. Moreover, the kernels can be re-used if illumination and projector conditions stay constant.
- In the example presented by Yamazoe (2008), eigenvalues rapidly decrease for high-order eigenfunctions. This property makes this formulation eligible for truncation. That is, omitting high-order eigenfunctions, and thus reducing the number of time-expensive Fourier transforms, leads to only a marginal loss of accuracy.
- In contrast to the sum of coherent systems (SOCS) approach (see next section), incident-angle-dependent mask spectra can be taken into account in a straight-forward manner. However, modifications of the mask also require the Fourier transforms to be recomputed.

3.2.6 Hopkins' Approach

An alternative reformulation of the partially coherent imaging equation (3.2) is proposed by Hopkins (1953). The main difference lies in an interchange of the order of integration, whose purpose is a separation of the influence of the object—the photomask, in the lithography case—and the imaging system including the pupil function and the illumination. The latter is often referred to as transmission cross-coefficient (TCC); it can be considered as a weight factor of the interaction between the two object frequencies (f', g') and (f'', g'') . The Hopkins formulation can be stated as follows:¹⁴

$$I(x_1, y_1) = \iiint_{-\infty}^{\infty} \mathcal{T}((f', g'), (f'', g'')) \mathcal{F}(O)(f', g') \mathcal{F}(O)^*(f'', g'') \exp(-i2\pi((f' - f'')x_1 + (g' - g'')y_1)) df' dg' df'' dg'', \quad (3.10)$$

where \mathcal{T} is the TCC given by:

$$\mathcal{T}((f', g'), (f'', g'')) := \iint_{-\infty}^{\infty} \mathcal{F}(J)(f, g) \mathcal{F}(H)(f + f', g + g') \mathcal{F}(H)^*(f + f'', g + g'') df dg. \quad (3.11)$$

¹³A similar approach using a different decomposition technique (termed principal component analysis) was proposed by Tsai et al. (2009).

¹⁴A detailed derivation is given by Born and Wolf (1999, pp. 603 and 604) and in the original communication by Hopkins (1953).

The TCC can also be thought of as four-dimensional linear filter resembling the response function (C.2) in the coherent imaging case (Wong, 2005, chap. 4, p. 70). Figure 3.7 illustrates this property: Each object point is band-limited by the pupil function. This is represented by the circles plotted around the spatial frequency points (f', g') and (f'', g'') —assuming a circular aperture. An interaction of these points occurs only if they have a common (overlapping) frequency domain. Moreover, only where the support of the Köhler illumination, that is, the mutual intensity distribution $(\mathcal{F}(J))$ coincides with this overlap, the TCC is non-zero. In other words, the intersection of the frequency domain of the two points and the illumination support form the integration domain of the transmission cross-coefficient.




Figure 3.7: Transmission cross-coefficient. The support of the TCC is given by the intersection of the effective source (J_C) and the frequency domain of the two object points under consideration $[(f', g')$ and $(f'', g'')]$. The frequency domain of the points is band-limited by the pupil, indicated by circles centered on the points.

For numerical convenience, the Hopkins formulation is often stated in terms of the spatial spectrum of the intensity in the image plane:

$$\mathcal{F}(I)(f, g) = \iint_{-\infty}^{\infty} \mathcal{T}((f' + f, g' + g), (f', g')) \mathcal{F}(O)(f' + f, g' + g) \mathcal{F}(O)^*(f', g') df' dg'. \quad (3.12)$$

The right-hand side of (3.12) exhibits a double integral as opposed to the quadruple one of (3.10) and thus drastically decreases the computational effort. On a discrete domain, the intensity spectrum $\mathcal{F}(I)$ can be efficiently transformed back into intensity values by means of a fast Fourier transform (FFT).

The TCC formulation allows for a feasible and effective way to analyze the imaging modulation of a partially coherent system and thus, for example, facilitates the selection of ideal filling factors for a given mask feature set (cf. Section 2.2.6). Besides, the separation of the object and the illumination/imaging system can be utilized to precompute and store TCCs and thus to further increase the computational efficiency of image calculations.

When applied for the design of circuits and photomasks, image computations have to be carried out on large (mask) areas. Computational efficiency plays hence a paramount role. For example, the comprehensive introduction of OPC software can be associated with the development of accelerated image models through approximations of the Hopkins formulation. Similarly to the Abbe approach, the partially coherent problem is decomposed into an incoherent sum of coherent images. In contrast to Abbe's method, however, the decomposition is not performed with respect to the source, but instead the Hopkins equation (3.10) is decomposed into a series of coherent kernels. A so-constructed infinite sum can in fact be shown to yield the same final image as the original Hopkins approach (Saleh and Rabbani, 1982). A finite, truncated sum can consequently be said to be an approximation whose error is to be minimized. A coherent decomposition based on an eigendecomposition was first proposed by

(Gamo, 1963) and later applied to lithographic imaging models by Pati and Kailath (1994). While the former authors' technique is based on the spatial-domain formulation shown in Equation (3.1), Cobb (1998, pp. 51–60) provides a derivation based on Hopkins' TCC formalism,¹⁵ making it amenable to traditional lithography simulation approaches. It can be summarized as follows: Let the image plane, the illuminator and the pupil function be represented by a discretization. The TCCs have a finite support, whose boundary can be given in terms of a maximum frequency (cf. Figure 3.7). Then the TCCs can be represented as Hermitian matrices.¹⁶ There exists an eigendecomposition—which can be numerical-efficiently conducted using an SVD—such that the TCC spectrum matrix can be written as:

$$\mathcal{F}(\mathbf{T})[(k_f, k_g), (l_f, l_g)] = \sum_m \lambda_m \mathbf{x}_m[k_f, k_g] \mathbf{x}_m^*[l_f, l_g], \quad (3.13)$$

where \mathbf{x}_m are the eigenvectors and λ_m are the eigenvalues of the TCC spectrum matrix. The indices (k, l) of the matrix $\mathcal{F}(\mathbf{T})$ correspond to discrete frequencies. The eigenvectors represent the coherent kernels of the system, the eigenvalues their weight. This becomes apparent when substituting the TCCs in the discrete version of the frequency form of Hopkins' approach [cf. (3.12)]:

$$\begin{aligned} \mathcal{F}(\mathbf{I})[k_f, k_g] &= \sum_{l_f} \sum_{l_g} \sum_m \lambda_m \mathbf{x}_m[k_f + l_f, k_g + l_g] \mathbf{x}_m^*[l_f, l_g] \\ &\quad \mathcal{F}(\mathbf{O})[k_f + l_f, k_g + l_g] \mathcal{F}(\mathbf{O})^*[l_f, l_g]. \end{aligned}$$

This equation can be re-ordered such that the sum over m becomes the outer summation (Cobb, 1998, pp. 54–55). Afterward the term of the then inner sums over frequencies is restated as a discrete convolution. By Fourier-back-transforming we then obtain the SOCS formula:

$$I[l_x, l_y] = \sum_m \lambda_m \left| (\mathcal{F}^{-1}(\mathbf{x}_m) * \mathbf{O}) [l_x, l_y] \right|^2. \quad (3.14)$$

Additionally both Pati et al. (1997) and Cobb (1998) propose to decompose the photomask into building blocks, i.e., basis functions that can be linearly combined into the mask pattern. Due to the translation-invariance and the linearity of convolution, each convolution term of (3.14) can be given as a linear combination of convolutions of the respective kernel and the mask basis functions. Furthermore, these individual convolutions can be precomputed and stored, commuting the intensity computation to a look-up/sum routine. This procedure also allows to locally update the image and hence to efficiently account for localized modifications of the mask that are, for example, due to OPC operations.

The decoupling of the illumination and projection system from the mask function that renders the Hopkins approach well suited for numerically efficient image computations necessarily constitutes a significant simplification: The field of the mask is assumed to be independent of the incident angle. For moderate NAs and feature pitches, especially when using a Kirchhoff mask model, only a marginal error results (Pistor et al., 2000b). However, state-of-the-art immersion lithography and mask types such as phase-shifting masks lead to a pronounced inclination-dependence (Erdmann et al., 2006a). As was shown previously (Section 3.2.5), the Abbe approach—also used in this work—presents a qualified method to include source-point-specific mask functions. Nonetheless, computation times requirements of (full-chip) OPC applications typically inhibit an exact treatment as provided by Abbe's method. Instead two type of approximations can be utilized:

- (1) EMF and incident-angle-specific correction terms are included into the Hopkins approach, which is then regarded as a compact model (cf., e.g., Liu et al., 2007). The correction terms have to be determined empirically, for example, through experiments or rigorous simulations.
- (2) As an alternative, a hybrid Abbe/Hopkins approach can be applied. Recently Adam et al. (2009) have proposed a method in which the illumination source is partitioned. For each partition, the mask spectrum is evaluated individually, for example, using a rigorous EMF solution. The intensity

¹⁵The relation between the two alternative approaches is shown by von Büнау (1997).

¹⁶In the two-dimensional case, the TCC depends on four parameters. Thus an appropriate mapping (stacking) operator from the four-dimensional matrix to a matrix must be defined.

contribution of a single partition is then computed using Hopkins' approach. Similarly to the Abbe method, the total intensity is obtained as an incoherent sum of the contributions of the individual partitions. The main assumption of this routine is that the EMF is approximately constant under small incident angle changes.

Although the focus of this work is on the improvement of lithographic process conditions and thus on small-area simulations that were conducted using Abbe's method, an application of the proposed optimization strategies to larger areas and using Hopkins' or a hybrid method will be discussed in the conclusions (Chapter 9).

3.3 Wafer Stack Light Propagation

The wafer stack consists of multiple layers. In addition to the photoresist, often anti-reflective coatings either below [bottom anti-reflective coating (BARC)] or, on top of the resist [top anti-reflective coating (TARC)] are used. While the former is mainly aimed at minimizing standing wave pattern that are due to self-interference with the back-reflected light, the latter is applied in order to reduce the impact of the resist thickness on the line width (often termed swing-curves).¹⁷ Therefore to compute the intensity distribution in the resist, the light propagation in all layers of the wafer stack has to be taken into account.

In the scope of non-planar wafer topographies, for example, for pattern doubling applications, a rigorous EMF computation has to be carried out. For each source point, the respective coupling spectrum has to be evaluated using a vectorial image computation. These spectra (or fields) are then supplied to the EMF solver [e.g., the waveguide method (or FDTD, respectively)], which in addition is provided with a representation of the wafer topography (e.g., see Shao et al., 2009).

For planar wafer topographies, less numerically complex approaches such as the *scaled defocus method* are applicable. There the focal plane is shifted according to the vertical position in the resist and the refractive indices of the stratified medium. Reflection is disjointly treated in the same manner (Bernard, 1988). Erdmann et al. (2006a) proposes a more rigorous approach using a *beam propagation method*, which is based on a finite-difference solution of the Helmholtz equation (A.15). Due to its formulation of boundary conditions, this approach can handle only small lateral variations of the material properties and offers thus a rather limited applicability to non-planar topography cases. Moreover, its time-complexity is only marginally lower compared to fully rigorous approaches using the waveguide method or FDTD.

In contrast, the transfer matrix method (TMM) is a numerically efficient formalism for the computation of the transmittance and reflectance in stratified media.¹⁸ It is an extension of Fresnel's equations. Contrariwise, it can also be considered a special case of RCWA approaches like the waveguide method: Each individual layer is assumed to be homogeneous in all directions. Thus in contrast to a general RCWA scenario, only the zeroth mode has to be evaluated, yielding a two-dimensional problem. In addition, the coupling of layers is significantly less numerically complex.

In order to apply the TMM to evaluate the intensity distribution in the resist, the wafer stack is represented by a stratified medium, consisting of homogeneous layers such as the incidence medium (i.e., air or water), the resist layer, an anti-reflection layer and a substrate, as shown in Figure 3.8.

Given the transmission coefficients for the entire stack (t) and for the sub-stack (t'), the transmitted amplitude of the stack (E_t) can be expressed both in terms of the incident wave in the zeroth layer (E_i) and in the resist layer (E_i^k):

$$E_t = E_i t = E_i^k t'.$$

Thus

$$E_i^k = \frac{t}{t'} E_i.$$

¹⁷The anti-reflectance on the resist/air interface leads to a relatively constant amount of light being coupled into the resist for different thickness values.

¹⁸A derivation of this method is given in Appendix D.




Figure 3.8: Wafer stack schematic: The goal is to find the electric field at any z position in the photoresist. The transmittance of both the entire stack (t) and the sub-stack below the resist (t') are given by the TMM. Furthermore, the reflectance of the sub-stack (r') and the electric field incident on the stack are known (E_i). This allows to determine both the incident field (E_i^k) and the reflected field (E_r^k) in the resist layer, which in turn can be used to evaluate the field at any z position in the resist, as explained in the text.

For the reflected field (E_r^k), we then have

$$E_r^k = E_i^k r' = \frac{t}{t'} E_i r'.$$

Let

$$k_0 h_{k,z} := k_0 n_k z \cos \theta_k$$

be the phase shift, similar to (D.5), where z is the relative position in the resist. The electric field in the resist can then be reformulated in terms of the wave functions:

$$E^k(z) = \frac{t}{t'} E_i e^{-ik_0 h_{k,z}} + \frac{t}{t'} E_i r' e^{ik_0 h_{k,z}}. \quad (3.15)$$

In the TE case,¹⁹ only the x and y components of the electric field are modified through the stratified medium; z contributions will only arise from the TM state. To conveniently integrate the intensity computation of the resist into the image formation formalism, Yeung et al. (1993) and Flagello (1993) propose a matrix notations that allows for a composition of the resist intensity distribution and the polarization correction from the exit pupil to the image plane (as explained in Section 3.2.2). The resulting flow can be schematically summarized as in Algorithm 3.2.

Algorithm 3.2: Resist image computation.

```

1 set number of  $z$  layers in resist;
2 foreach source point do
3   foreach spatial frequency pair ( $f_x, f_y$ ) do
4     determine incidence angle on stratified medium;
5     retrieve far-spectrum from imaging algorithm;
6     evaluate transmittance and reflectivity using TMM;
7     foreach  $z$  layer do
8       retrieve  $E_x$  and  $E_y$  contribution arising from TE mode (3.15);
9       retrieve  $E_x$ ,  $E_y$  and  $E_z$  contribution arising from TM mode;
10      perform polarization correction (exit pupil to image plane);
11   foreach  $z$  layer do
12     transform spectrum to field;
13 foreach  $z$  layer do
14   superpose fields;
```

During exposure, the resist properties undergo a modification known as bleaching. Although for thin resist types and especially CARs that are typically applied in state-of-the-art projection lithography, accounting for bleaching within the exposure kinetics models is often considered to be sufficient, a rigorous treatment, and an according intensity distribution correction has also been proposed (Yeung, 1985; Bernard, 1988; Erdmann et al., 2001).

3.4 Resist Modeling

Besides the optical system and its components, the photo-chemical reactions that the photoresist undergoes in each process step are a major domain of computational lithography. Due to the relative high numerical effort required for a rigorous simulation of the chemical resist properties, most resolution enhancement optimization routines employ only simplified models. In this work as well, a simple threshold model instead of a resist simulation is applied. In order to illustrate the relation between the simplification and the full simulation, and for the sake of completeness, the main considerations of resist modeling are briefly summarized.

¹⁹where the electric field vectors are perpendicular to the plane of incidence, compare Figure D.1

3.4.1 Pre-bake

Evaporation and diffusion kinetics during the pre-bake step can be given as follows:

$$\frac{\partial M(z)}{\partial t} = -kM + \nabla(D_M \nabla M),$$

where M is the solvent concentration, k is the evaporation rate constant, D_M is the solvent diffusion coefficient. While by neglecting diffusion or assuming Fickian diffusion an analytic solution can be given (Mack, 1998), non-Fickian diffusion requires the application of numerical (most prominently, finite-difference) approaches (Byers et al., 2003).

3.4.2 Exposure Kinetics

A common model for the description of the exposure kinetics was introduced by Dill et al. (1975) and is hence often referred to as Dill's model. Its derivation can be briefly summarized as follows. According to Lambert's phenomenological law, the (time-dependent) intensity distribution in an absorbing medium is given by

$$\frac{\partial I(z; t)}{\partial z} = -\alpha(z; t)I(z; t), \quad (3.16)$$

where $I(z; t)$ is the time-dependent intensity at a given position z , and α is the time- and position-dependent absorption coefficient. Beer's law states that for dissolute solutions (absorbing species in a non-absorbing fluid), the intensity in the resist is proportionate to the concentrations C_i of the absorbing species, whose absorption coefficients are denoted by a_i :

$$\alpha(z; t) = \sum_i a_i C_i(z; t). \quad (3.17)$$

The typical photoresist species are given as follows:²⁰

- (1) M : *photo-active compound (PAC)*, which in the case of CARs is a PAG. M_0 is the initial concentration prior to exposure.
- (2) P : *exposure products* due to reaction: $M \xrightarrow{\text{UV}} P : P = M_0 - M$.
- (3) S : *solvent(s)*.
- (4) R : *resin*.

Yielding the combined absorption coefficient

$$\alpha = \sum_{\mathfrak{S} \in \{M, P, S, R\}} a_{\mathfrak{S}} \mathfrak{S}. \quad (3.18)$$

In order to separate (3.18) into a bleachable component (A) and a non-bleachable one (B), Dill proposes the following reformulation:

$$\alpha = Am + B, \quad (3.19)$$

where

$$\begin{aligned} A &:= (a_M - a_P)M_0 && \text{bleachable Dill } A \text{ parameter} \\ B &:= a_P M_0 a_R R + a_S S && \text{non-bleachable Dill } B \text{ parameter} \\ m &:= M/M_0 && \text{relative PAC concentration.} \end{aligned}$$

In addition to numerical resist modeling, the Dill parameters play a central role for the determination of resist properties as they allow for a simple measurement set-up in that only the fully exposed and the unexposed resist absorption coefficients need to be found. The third and last Dill parameter (C) is given by the rate equation:

$$\frac{\partial m}{\partial t} = -CI m, \quad (3.20)$$

where m is defined as before and I is the absorbed intensity. This is to say that the PAC dissolution rate is related to the absorbed intensity and time-dependent availability of PAC (Kirchauer, 1998).

²⁰From here on, the locus- and time-dependencies are omitted for readability.

3.4.3 Delay

The time that elapses between exposure and subsequent process steps, specifically post-exposure bake, is often referred to as delay time. During this inevitable delay, CARs are known to suffer from a premature acid loss, for example, as a result of evaporation on the resist surface and, more importantly, through neutralization of the acid and base contaminants. While typical models ignore the former effect, the latter, the acid/base neutralization, can be accounted for by approximating the alkaline contamination concentration as follows (Mack, 1998, 2007):

$$B(z, t) = B_0 \operatorname{erfc} \left(\frac{z}{\sqrt{2}\sigma} \right), \quad (3.21)$$

where σ is the diffusion length (assuming constant diffusivity), z is the depth coordinate and B_0 is the base concentration on top of the resist. It is assumed to be constant. The complementary Gaussian error function (erfc) is defined as:

$$\operatorname{erfc} : x \mapsto 1 - \operatorname{erf}(x) = \frac{2}{\sqrt{\pi}} \int_x^\infty e^{-t^2} dt.$$

Under the presumption that the entire base reacts with the typically strong acid, the acid concentration directly before the post-exposure bake step (A') can be approximated by:

$$A' = A - B,$$

where A is the acid concentration directly after exposure and B is the base concentration as in (3.21). A detailed derivation of diffusive acid loss both during the delay time and in the subsequent post-exposure step is also presented by Weiß et al. (1995).

3.4.4 Post-Exposure Bake (PEB)

When applied to CARs, the post-exposure bake induces the deprotection of polymers, catalyzed through the acid generated during exposure. In addition, diffusion has to be taken into account and coupled with the deblocking reaction.

In summary, the following processes have to be modeled (cf. Schnattinger, 2007, pp. 42–44, and Erdmann, 2010):²¹

- (1) Acid/base neutralization: $A + B \longrightarrow X_{AB}$,
- (2) Acid-catalyzed deprotection: $A + M \xrightarrow{\Delta T} X_M$,
- (3) Non-catalyzed deprotection: $M \xrightarrow{\Delta T} X_M$,
- (4) Acid decomposition: $A \xrightarrow{\Delta T} X_A$,
- (5) Base decomposition: $B \xrightarrow{\Delta T} X_B$,
- (6) Acid diffusion,
- (7) Base diffusion,
- (8) Indiffusion of alien species, and
- (9) Outdiffusion of species,

where A and B are the base concentrations, respectively, M is the concentration of the protected polymer sites, and $X_{\mathfrak{S}}$ ($\mathfrak{S} \in \{A, B, AB, M\}$) denotes the respective reaction product. Thermally-induced reactions are denoted by $\xrightarrow{\Delta T}$.

Several authors proposed compact formalisms for the coupled reaction and kinetic processes above (see, for example, Capodiec et al., 1995; Henke and Torkler, 1999), featuring concentration-dependent diffusion terms. In addition to standard Fickian diffusion, they allow thus for other diffusion models

²¹The base referred to here is typically called quencher base. Its main purpose is to improve contrast, for example, by quenching acid in unexposed resist areas, where stray light leads to an undesired acid generation.

including linear and exponential diffusion. The following highly flexible partial differential equations (PDEs) formalism was proposed by Henke and Torkler (1999):

$$\frac{\partial M}{\partial t} = -k_1 M^p A^q - k_2 M \quad (3.22a)$$

$$\frac{\partial A}{\partial t} = -k_3 A^r + \nabla(D_A \nabla A) - k_4 AB \quad (3.22b)$$

$$\frac{\partial B}{\partial t} = -k_4 AB + \nabla(D_B \nabla B) - k_5 B, \quad (3.22c)$$

where D_A and D_B denote the diffusion coefficients of acid and base, respectively. The constant rate factors are given by k_i , and the species concentrations notation is as before. Equation (3.22a) represents the deprotection reaction. The parameter k_1 denotes the constant deprotection rate; k_2 is the constant rate of additional thermally induced deprotection, which is not fully avertable. Parameters p and q specify the reaction order of the respective species. The acid loss is given by (3.22b); r is the reaction order and k_3 is the constant acid loss rate. As Henke remarks, the base-acid neutralization described in (3.22c) could also be included in the acid loss term (3.22b). It is introduced, however, to account for a potential non-stability of the base concentration profile over time. Again k_4 is the constant rate factor, this time for the base-acid neutralization. The rate for the base loss is denoted by k_5 . This model can be further generalized to include diffusion terms that may depend on combinations of concentrations of several species (Matiut et al., 2003).

The coupled PDEs (3.22a)–(3.22c), with the initial condition being the acid concentration after exposure, can be numerically solved using a finite-difference technique. Different boundary conditions can be applied, for example, using periodic boundary conditions (examples of numerical schemes are given by Li, 2001; Fuhrmann et al., 2010).

3.4.5 Develop

In order to obtain the final resist profile, the topography change due to the development has to be computed and represented appropriately. For that purpose, the development rate is evaluated that is then used in a geometric representation of the resist.

All models employed to compute the rate of coupled reaction and diffusion processes of lithographic processes are of semi-empirical nature.²² They depend on a number of parameters, either physical or non-physical, that need to be determined externally, i.e., through experimental measurements. An entirely empirical model was proposed by Dill et al. (1975). Their model depends on three (non-physical) parameters that are determined through a fit with experimental data. Kim et al. (1984) developed a similar approach in which the dependence to fit parameters is attenuated by introducing two measurable parameters (the rates in the unexposed and in the fully exposed resist case). In that model, another effect, known as *surface inhibition*, is also accounted for: In the vicinity of the resist surface, there is a significant decline of the dissolution rate. The exact reason for this phenomenon is still subject to research Burns et al. (2001). In contrast to Kim et al., who describe surface inhibition as a z -dependent (exponential) factor of the development rate, Weiß et al. (1995) proposes the concept of an *effective acid concentration*. There the z -dependent function is a factor of the acid rather than the development rate. The actual development rate in the Weiß model is then a function of the effective acid concentration and measured parameters obtained with a development rate monitor (DRM).

Earlier Mack (1987) developed a slightly different development rate formulation, today known as the Mack model. It is based on the following mechanisms:

- (1) Diffusion of the developer to the resist surface. The rate (r_D) is given by:

$$r_D = k_D(D - D_s), \quad (3.23)$$

where D is the developer concentration in the resist bulk, D_s is the same concentration at the surface of the resist and k_D is the rate constant.

²²For example, in that they constitute numerical models, whose parameters are obtained through significant experimentation.

- (2) The reaction of the developer with the resist. If P , as before, is the concentration of deprotected inhibitors, n is the number of inhibitor sites that protect the resin polymer from dissolution, and k_R is the rate constant, the reaction rate r_R can be described by:

$$r_R = k_R D_S P^n. \quad (3.24)$$

- (3) Diffusion of the dissolved resist back into the solution. This process is assumed to be approximately instantaneous; thus no rate function is employed.

In addition to these considerations, Mack argues that since the reactions, (3.23) and (3.24), take place serially, they result in a steady state, i.e., $r_R = r_D =: r$. By equating the former formulae, one can solve for D_S , yielding the following original Mack model:

$$r = r_{\max} \frac{(a+1)(1-m)^n}{a + (1-m)^n} + r_{\min} \quad (3.25)$$

Here m is defined as the relative inhibitor concentration [see (3.19)], and the parameter a is set to be constant with:

$$a := \frac{M_0^n k_D}{k_R}.$$

The value of the maximum rate r_{\max} is given by:

$$r_{\max} = \frac{D k_D}{a+1}.$$

Finally the parameter r_{\min} is the dissolution rate of unexposed resist, which has been experimentally shown to be unequal to zero.

Several extensions for the specifics of CARs, to account for high resolution resists and to fit to facilitate calibrations to experimentally obtained data were proposed, leading to an *enhanced Mack model* (Mack, 1992), the *notch model* (Mack and Arthur, 1999) and an *enhanced notch model* (Robertson et al., 2001).

With the local development rates, obtained from any of the above models, the resist profile can be evaluated. Due to their numerical stability, two methods are predominantly used in state-of-the-art lithography simulators:

- (1) The *cell removal* (Henke et al., 1991; Hirai et al., 1991) method can be regarded as a cellular automaton. The resist layer is represented by a three-dimensional equidistant cubic grid, in which each cell is characterized by its rate, depending on the previous process steps (specifically exposure and post-exposure bake) and the corresponding local development rate. In addition, each cell is assigned a counter to with an initial value of v , representing its volume. Cells with a volume of zero are treated as dissolved.²³ The total develop time is partitioned into discrete time intervals. With each time step, all cell counters are modified simultaneously, using the following update scheme: If a cell that represents an exposed portion of the resist has neighbors representing developer, its counter is decremented according to its local dissolution rate and the number of adjacent developer cells (and, of course, the duration of the time interval). As a result, the final resist profile is given after the last time step by the non-dissolved cells.
- (2) The *level set method* (Osher and Sethian, 1988) takes a significantly different approach, using an implicit representation of the advancement of fronts—or similarly, the regression of resist due to dissolution (Sethian and Adalsteinsson, 1997). Given a family of simple closed curves $\gamma(t)$ in \mathbb{R}^n , where t is the time at which the curve is to be considered. Additionally, a function v , yielding the velocity in the normal direction of each point on γ , is to be defined. This function may depend on local, global and independent properties such as development rates. The level set function is then

²³Negative values are sometimes allowed and called *overkill* values. If a cell exhibits such a negative counter, the excessive “development force” is passed to adjacent cells.

a function that relates all points of interest \mathbf{x} to its distance d to the curve (or more precisely, a hypersurface) at a given time:

$$\begin{aligned} \phi : (\mathbb{R}^n; t) &\rightarrow \mathbb{R} \\ \phi(\mathbf{x}; t) &= \begin{cases} +d & \text{if } \mathbf{x} \text{ is inside of } \gamma(t) \\ 0 & \text{if } \mathbf{x} \text{ is on } \gamma(t) \text{ (the zero level) ,} \\ -d & \text{if } \mathbf{x} \text{ is outside the } \gamma(t) \end{cases} \end{aligned} \quad (3.26)$$

The initial condition is given by laying down the level set function at $t = 0$, for which the initial front, for example given by deprotected polymer sites, is known. With this initial condition, the velocity function v and the time derivative

$$\frac{\partial \phi}{\partial t} v |\nabla \phi| = 0, \quad (3.27)$$

we obtain a (Hamilton-Jacobi) differential equation of dimension one higher than the original problem. For practical reasons, curves and level set functions are commonly defined on a discrete grid. This approach also allows for the application of numerical approximations such as finite differences. In the lithography development case, resist is dissolved; thus the front only moves in the negative direction. This property leads to significant speed-up as it allows for a formulation using a stationary-time eikonal equation and is solvable with an algorithm called *fast marching method* (Sethian, 1996).

As indicated, both methods are typically carried out on a Cartesian grid. In order to use the obtained topological data from either method, a geometrical post-process step is often required. The so obtained profile can be used to conduct metrology operations similar to the aerial image evaluation described before, but under the influence of the photochemical resist kinetics. In addition, measurements of three-dimensional properties such as side-wall angles or fluctuations such as linewidth roughness (LWR)—global variations of the CD—and line-edge roughness (LER)—local, often high frequency, variation at one edge—can be performed.²⁴

For a more rigorous treatment of LER and LWR and especially to account for the stochastic nature of reaction and diffusion processes, mesoscopic resist models have been proposed (Mülders et al., 2005; Schnattinger, 2007; Biafore et al., 2009). There, concentrations and rates of the continuous formulations have to be represented as probability chains. Stochastic resist models are especially useful to study the electron trajectories observed in NGL options (see, for example, Kotera et al., 2008).

Both continuous and stochastic resist models depend on a multitude of physical and mathematical parameters that need to be adapted in order to closely resemble real conditions. While some of these parameters can be measured, others can only be determined through calibrations. For example, a calibration technique based on evolutionary algorithms has been proposed elsewhere (Tollkühn et al., 2003). Rigorous resist models are numerically expensive and hence time-consuming. As an alternative, non-physical, simplified models are often applied whenever fast computations are a concern, such as in OPC or source/mask optimization. Especially in light of the calibration efforts required to fit parameters of “physical” models to real-world situations and the resulted uncertainties, well-calibrated simplistic models are not less predictive per se. Of course, care has to be taken as to their scope of validity. Examples of simplified resist models include a *lumped resist model* (Mack, 1994)—a three-dimensional extension was later proposed by Byers et al. (2002), who later added a CAR model (Byers et al., 2004)—, the *diffused aerial image model [DAIM]* (Ahn et al., 1998) and the *acid dose diffusion threshold model [ADDIT]* (Van Steenwinckel and Lammers, 2003).²⁵

As we shall see in Section 7.2.2, in this work, a further simplification, an aerial-image-based threshold model, is employed. This technique is often applied to save additional computation time and for non-resist-specific OPC or optimization approaches (cf., e.g., Cobb, 1998). An extensive study comparing the errors of different threshold-based models was conducted by Fuard et al. (2002).

²⁴There is no concise definition for LWR and LER. Here, we follow Mack (2007).

²⁵A summary and a number of extensions to these models are given by Tollkühn et al. (2004).

4 Genetic Algorithms

An optimization goal aiming at improving the lithographic process can be considered to be fairly general and hence arbitrary. In fact, such an abstract problem definition opens up a vast field of diverse sub-problems all of which may exhibit entirely different properties. But even without considering a specific problem, the previous chapters already suggest that any task in this setting will involve a multitude of physical and chemical aspects, leading to a large number of both optimization variables and criteria, with nonlinear interrelations. The complexity is further increased under the assumption that in general search spaces will be neither continuous nor convex and additionally feature multiple local extrema and incommensurable figures of merit. One could argue that because of these expected properties no optimization algorithm will prove suitable. And indeed, a general approach is not foreseeable. On the other hand, the actual formulation of the problem presents an important degree of freedom and ultimately determines the class of the problem and in turn whether it is amenable to a specific optimization approach or not. Examples on how to adapt lithographic problems to numerical optimizers are given in Section 7.1.

In this work, we have investigated the application of a meta-heuristic, population-based approach, termed *evolutionary algorithm (EA)* to several lithographic optimization problems. In contrast to conventional numerical optimizers, these types of algorithms do not exploit any search-function-specific information such as derivatives. Rather the search space is adaptively sampled as the algorithm proceeds. Improvement is achieved solely by an evaluation and comparison of the search function values. Thus the only requirement is the ability to obtain the function value for any parameter configuration that lays within the predefined optimization variable domain.

Evolutionary algorithms have become very popular over the last three decades, although the basic idea, the utilization of mechanisms observed in the process of natural evolution, has been introduced far earlier, in the late 60s. A multitude of different variants have been proposed, all of which use operators such as selection, recombination and mutation, mimicking their natural role model (Darwin, 1998). After a brief outline of related work, this chapter will focus on the critical aspects of *genetic algorithms*, which were used in this work. The last section discusses the application of genetic algorithms to multi-objective optimization problem formulations.

4.1 Related Approaches

Metaheuristic search approaches can be coarsely subdivided into *single solution* and *population-based* techniques. A popular example for the former is the simulated annealing (SA) approach (Kirkpatrick et al., 1983), in which the states (candidate solutions) are distributed according to the Boltzmann distribution, constituting a thermal equilibrium at a certain temperature. By decreasing the temperature, the distribution changes so that it strongly favors low energy states. This temperature decrease is to be performed slowly in order for the distribution to be accurately and precisely sampled.

Another metaheuristic approach, specifically aimed at solving combinatorial problems, is called tabu search (TS) (Glover, 1986). Broadly speaking, this method can be considered a globalized descent method through dynamic adaptation of (local) feasibility conditions. The rationale behind this procedure is to allow the optimizer to continue its search even if a local optimum has been found, which conventional descent methods would yield as a final solution. Permitting the optimizer to proceed further, however, easily leads to a cyclic behavior, with local optima constantly being revisited. In tabu search, this is prevented by storing previously visited candidates (or moves) into a list, called tabu list: New solutions in the neighborhood of the current solutions are valid only if they are feasible with respect to this tabu list. For efficiency reasons, it is common practice to store the moves that lead to

candidate solutions rather than the actual visited candidates. This may, however, lead to an unintentional tabuing of admissible solutions. To relax this so-caused exceedingly restrictive behavior, tabu search also maintains a list of aspiration conditions, for example, allowing for globally best observed candidates to be visited despite their not being feasible according to the tabu list. Another feature of TS is its introduction of two additional search stages, dedicated to an *intensification* of the local search and a *diversification*, a wider spread, of solutions. This is achieved by modifying the objective function to reward new candidates that are in close vicinity to the current solution or far off, respectively. Commonly, each of these two stages is executed alternately with the third, the regular search phase, in which the unaltered fitness function is employed. Each stage and the entire flow can be bound to either standard—maximum number of iterations or an insufficient improvement of solutions—or tailored termination criteria. Tabu search has been successfully applied, especially in the field of routing and scheduling problems (e.g. Hertz, 1991).

Another heuristic approach is the down-hill simplex method proposed by Nelder and Mead (1965). It heuristically exploits local zero-order information of the search space. Since we have used it in the scope of a hybrid global/local search approach, more details will be presented later in Section 5.4.2.

The remainder of this section is devoted to population-based techniques, often subsumed under the term *artificial evolution*.

4.1.1 Evolution Strategy

One of the pioneering works in the field of artificial evolution with a strong emphasis on engineering applications was conducted by Rechenberg (1973) and later implemented and extended by Schwefel (1977). In evolution strategy (ES) approaches, optimization variables are typically represented as real numbers. Each of these variables, often referred to as objective variables \mathbf{x} , is supplemented by a strategy parameter \mathbf{s} , which can be interpreted as an internal parameter that determines the interaction between the current variable value and the environment—for example, the step length or standard deviation of modifications in subsequent iteration steps. This strategy parameter may consist of a single element, which is applied to all object parameter components, or a vector of the same rank, allowing each object parameter to have a dedicated strategy counterpart. Formally, for the latter case, we have that,

$$\mathbf{x}, \mathbf{s} \in \mathbb{R}^n.$$

Similarly to conventional numerical optimizers, the objective vector \mathbf{x} serve as the argument for a merit function ($f : \mathbb{R}^n \rightarrow \mathbb{R}$) to be minimized (or maximized):

$$\begin{aligned} & \arg \min_{\mathbf{x}} f(\mathbf{x}) \\ & \text{subject to } x_i^l \leq x_i \leq x_i^u \quad (i = 1, \dots, n), \end{aligned}$$

x_i^l and x_i^u denote optional lower and upper variable bounds, respectively. ESs are iterative approaches. That is, candidate solutions (individuals) are iteratively modified in order to achieve improvements that eventually lead to a convergence toward the desired extremum. Since ESs are probabilistic approaches, thus lacking an exact measure for the accomplishment of an optimum, a termination criterion for the iterative procedure has to be predefined. This may include a maximum number of steps, a certain merit threshold or a number of iteration steps with no (significant) improvement.

Individuals are modified through mutation, a random, minor variation of the original candidate solution, which can for example be achieved by adding normally distributed random numbers, where the strategy parameters denote the standard deviation. Given the objective variables \mathbf{x} and the strategy parameters \mathbf{s} , the modified objective vector is obtained by,

$$\mathbf{x}' := \mathbf{x} + \mathcal{N}_0(\mathbf{s}),$$

where $\mathcal{N}_0(\mathbf{s})$ denotes a vector of normally distributed random numbers with mean 0 and standard deviations s_1, \dots, s_n ,¹ i.e.,

$$\mathcal{N}_0(\mathbf{s}) = [\mathcal{N}(0, s_1), \dots, \mathcal{N}(0, s_n)]^T.$$

¹In the single strategy parameter case, the \mathbf{s} can thought of, containing n identical values.

With the outlined approach, the standard deviation, which determines the mutation strength, has to be preset prior to optimization. Alternatively, the mutation strength can itself also be subjected to adaptation. Rechenberg (1973), for example, proposed what is known as “1/5th success rule:” After a number of generations, the success of the current mutation rates is evaluated in terms of the percentage of improved offspring solutions. If 1/5th of the offspring solutions excel their parents, the mutation rate is decreased; if 1/5th perform worse, it is increased; in all other cases, the rate is left unchanged. This scheme is typically applied to the scalar strategy parameter case, but can be extended to strategy vectors in a straightforward manner. Schwefel (1977) observed that this simple adaptation may lead to a premature convergence by a collapse of step sizes and a consequent reduction of the search space to only a subset of the original dimension. He therefore introduced an adaptation scheme in which the standard deviations are correlated such that the resulted distribution eventually follows the search direction (Schwefel, 1987). In other words, with this technique not only the step size but also the search direction is evolved according to (local) search space properties.

In its simplest form, termed $(1 + 1)$, a single individual is maintained and improved throughout the entire ES run. There, an update step is given as follows: First, the current individual, whose pay-off value can be assumed to have been previously evaluated, is cloned. The duplicate is then modified through a mutation operator, yielding two individuals, the original candidate solution (parent) and the mutated variant (child). Finally, the child’s pay-off is evaluated, and only the individual with the higher merit is kept. This scheme leads to a hill-climbing behavior quickly converging to local extrema. Moreover, since like any artificial evolution approach, ES does not evaluate local information such as gradients, the search direction implicitly resulting from the child individual may not be optimal in that it directs the subsequent steps toward a suboptimum. This situation is depicted in Figure 4.1(a), where a modified solution performs better than its parent, but worse than potential alternative ones. By generating not only one but multiple (λ) mutated variants, a broader search radius is obtained, increasing the exploration power. (Figure 4.1(b)). This scheme is called $(1 + \lambda)$. Another variant of this scheme $(1, \lambda)$ discards the parent from selection after mutation. This scheme can be used to further slacken convergence to a local optimum, which may occur in the presence of so-called super-individuals that exhibit a pay-off so high (in case of maximization) that they prevent offspring solutions to replace them—and in turn prohibit an exploration of other regions of the search space.

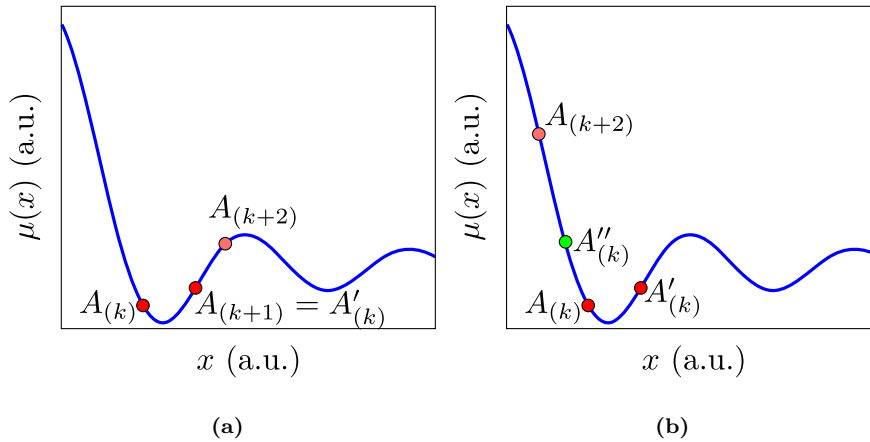


Figure 4.1: Evolution strategy convergence comparison of two selection scheme: (a) Using a $(1 + 1)$ selection scheme, that is, the best of the parent and one mutated child is selected, may lead to a suboptimal search direction. (b) By generating multiple offspring individuals chances are that one of them exhibits a merit that eventually directs convergence toward a better path. The iteration step is denoted by k , individuals (A) resulting from mutation are indicated by primes $'$, $''$.

To better adapt ES to multimodal search spaces, it can be transformed into a population-based approach by maintaining multiple individuals per iterative step. The obtained schemes are $(\mu + \lambda)$ and (μ, λ) . Out of a given population of μ evaluated individuals, λ individuals are randomly selected and mutated. The μ best solutions including $(\mu + \lambda)$ or excluding (μ, λ) their parents then form the new population. The size of both the population (μ) and the selection pool (λ) strongly depend on the problem. In general, they can only be determined empirically. Two of the most important actuating parameters are the number of optimization variables and the number of local optima in the search space.

Using a mutation-only approach, early ES variants can be considered hill-climbers with adaptive step lengths, applied in multiple regions in the $(\mu + \lambda)$ or the (μ, λ) cases. Similarly to genetic algorithms, explained later, ES can also be extended to incorporate recombination operations. The motivation for their application can be broadly summarized as follows: With proceeding iterations, individuals can be expected to possess parameter configurations that have been assessed superior compared with other solutions. Intermixing these solutions may eventually generate individuals that combine conducive parts and thus lead to an improvement. It should be noted that this procedure is a probabilistic one and hence depends on a large number of recombination operations.² Recombination in evolution strategies can be performed on an arbitrary number of individuals, that is, two, three or more individuals can be selected to mate. The recombination procedure can be classified according to the level of intermixing it provides, namely:

- (1) With *discrete recombination*, each parameter (both objective and strategy) takes the value of one of the parents, determined randomly. The following example demonstrates a simplistic case in which two individuals are recombined to form one offspring solution:

$$\begin{aligned} A_1 &: [[3.5, 2.7, -8.9]^T, [1.2, 2.2, 2.3]^T] \\ A_2 &: [[1.3, 3.3, -10.5]^T, [1.6, 2.4, 0.7]^T] \\ A' &: [[1.3, 2.7, -10.5]^T, [1.2, 2.2, 0.7]^T]. \end{aligned}$$

- (2) *Intermediate recombination* is a continuous recombination approach in which the average of the parent variable values yield the offspring. For example,

$$\begin{aligned} A_1 &: [[3.5, 2.7, -8.9]^T, [1.2, 2.2, 2.3]^T] \\ A_2 &: [[1.3, 3.3, -10.5]^T, [1.6, 2.4, 0.7]^T] \\ A' &: [[2.4, 3, -9.7]^T, [1.4, 2.3, 1.5]^T]. \end{aligned}$$

Many variants of this approach exist. Typically the average is weighted by a uniform random variable that determines the influence of either individual:

$$\mathbf{x}' := \mathbf{x}_1 + \mathbf{U} \circ (\mathbf{x}_1 - \mathbf{x}_2),$$

where \mathbf{x}_1 and \mathbf{x}_2 are the objective parameters of the parent solutions, \mathbf{U} is a vector containing uniformly distributed random numbers $[\mathbf{U} \in [0, 1] \text{ } (i = 1, \dots, n)]$ and \circ denotes a component-wise multiplication (Hadamard product).

Evolution strategies have been successfully applied to a variety of different problems, especially in the regime of numerical optimization.

4.1.2 Evolutionary Programming

Almost simultaneously with Rechenberg's and Schwefel's work on evolution strategies, a closely related approach—albeit aimed not at numerical optimization but at automatically evolving finite state machines to predict input sequences—was proposed by Fogel et al. (1966).³

²The role and benefits of recombination in evolution strategies and other evolutionary algorithms has been widely and controversially debated (for example, see Fogel and Atmar, 1990; Spears, 1998).

³Surprisingly enough, evolution strategy and evolution programming researchers first made contact in the early 1990s. The main similarities and differences between the two approaches have been first summarized by Bäck et al. (1993).

Like in ES, the decision space in evolution strategy (EP) is typically represented by real-numbered vectors, where each component corresponds to one optimization variable. After a random initial population with size μ has been generated and its candidates evaluated according to a fitness function, each parent produces one mutated offspring solution, yielding a selection pool double the size of the population (2μ). All new individuals in the selection pool are evaluated. To reduce the number of individuals to the original population size, selection is performed in the following way: In a first step, each solution receives a score representing its fitness. For that, the respective solution is compared with $q \geq 1$ individuals randomly drawn from the selection pool. The μ individuals with the highest score then constitute the next population. This process is repeated until the termination criterion is met. Originally using a static Gaussian mutation operator, controlled by user-provided strategy parameters, a great variety of dynamic, adaptive and self-adaptive mutation schemes have been later devised. There, strategy parameters similar to the ES approach are typically used to subject the mutation strength to evolution (Fogel et al., 1991). Similarly, due also to the mutual inspiration of the different artificial evolution communities, numerous alternative selection operators have been proposed.⁴

In contrast to other evolutionary algorithms, EP does not use recombination. Fogel et al. (1966) argued that candidate solutions resemble species representatives rather than individuals, and there is not inter-species information exchange through crossover. According to Fogel and Atmar (1990) an introduction of crossover even degraded the convergence behavior of EP. Their study was based on rather simplistic test functions. Because this observation strongly contradicting other findings, Schaffer and Eshelman (1991), for example, conjecture that the reason for the abysmal crossover performance in that study could also be attributed to the decision space representation.⁵ Later Fogel and Stayton (1994) extended the former analysis to a larger test bed and compared it with different GA versions, one specifically with a tailored representation for numerical optimizations (Schraudolph and Belew, 1992).⁶ The conclusion the authors drew from their experiments, stating that the success of evolutionary function optimization does not depend on recombination, can definitely be said to be correct (as for example later proved by the no free lunch (NFL) theorem of Wolpert and Macready, 1997). It does not, however, reveal any insight into the benefits or drawbacks of either crossover or the employed decision space representation.

Like other artificial evolution approaches, EP has a proven history of numerous successful applications in the field of combinatorial and function optimization approaches.

4.1.3 Differential Evolution

Differential evolution (DE) has been developed by Storn and Price (1997) and quickly gained attention and recognition as it proved to be highly competitive for a number of (numerical) test problems (Price, 1996). The flow of DE is illustrated in Algorithm 4.1: A fixed-size population is randomly initialized. In one iteration step, new candidate solutions are created for each slot of the population. First, an intermediate solution is obtained by adding a difference vector of two solutions, which can be scaled by a user-defined differential weight (F), and a third solution (l. 5). Thus a total of three individuals are combined, all of which are randomly selected such that they are mutually different and do not correspond to the former solution of the target slot in the population. Despite its involving a combination of candidates, the authors refer to this step as mutation. It is considered the most important operator in differential evolution. In their theoretical study on the convergence of DE, Dasgupta et al. (2009) suggest that this procedure can be regarded as a stochastic approximation of the gradient of the objective function. This hypothesis is also consistent with the original DE author's observation that a premature convergence toward a local extremum can be slowed by additionally employing a crossover operator (ll. 6), which recombines the intermediate solution obtained from mutation with the previous agent of the target population slot. This procedure is very similar to the discrete recombination of evolution strategies. In addition to uniformly randomly selecting values either from the intermediate or the former solution—whose proportion can be user-adjusted by a crossover probability parameter

⁴For a brief survey on state-of-the-art operators and variants of EP, see, for example, Engelbrecht (2007, pp. 189–204).

⁵A similar observation was reported by Bäck et al. (1993).

⁶See Section 4.3 for further details on parameter representation.

(CR)—at least one component has to originate in the newly create intermediate solution in order to prevent stagnation. The so-composed agent then replaces the former solution if and only if it exhibits a higher pay-off (ll. 11). In analogy to other evolutionary algorithms, this operation is commonly referred to as selection.

Algorithm 4.1: Basic Differential Evolution Algorithm

```

1  $k \leftarrow 0$  ; // iteration number
2 randomly initialize  $\mathbf{x}_i^{(0)} \in \mathbb{R}^n$  ( $i = 1, \dots, m$ );
3 while not termination criterion met do
4   for  $i \in [1, m]$  do in parallel
      // Mutation
5      $\mathbf{v} \leftarrow \mathbf{x}_{\sigma(1)}^{(k)} + F \left( \mathbf{x}_{\sigma(2)}^{(k)} - \mathbf{x}_{\sigma(3)}^{(k)} \right)$  ; /* where  $F$  is the user-provided differential
      weight---typically  $F \in [0, 2]$ ---and  $\sigma$  is a uniform random permutation of
       $\{j \mid 1 \leq j \leq m \wedge j \neq i\}$  */
      // Recombination
6      $r \leftarrow$  uniform random number with  $r \in [1, \dots, n]$ ;
7     for  $l \in [1, n]$  do /* all parameters (i.e., vector components) */
8       if  $\mathcal{U} \leq \text{CR} \vee r = l$  then /* uniformly drawn number  $\mathcal{U}$  smaller than crossover
      rate or random component hit */
9          $x_{i,l}^{(k+1)} \leftarrow v_l$ ;
10      else
11         $x_{i,l}^{(k+1)} \leftarrow x_{i,l}^{(k)}$ ;
      // Selection
12     if  $f(\mathbf{x}_i^{(k)}) > f(\mathbf{x}_i^{(k+1)})$  then
13       // Former solution better: reject new individual
14        $\mathbf{x}_i^{(k+1)} \leftarrow \mathbf{x}_i^{(k)}$ ;
15    $k \leftarrow k + 1$ ;

```

Many variants of this algorithm have been proposed, for which the original authors use the following notation:

$$\text{DE}/x/y/z :$$

x is the origin vector which is modified by the weighted difference. It can take values like “**rand**”, in which case this vector is randomly selected from the preceding population, as has been demonstrated in Algorithm 4.1. Alternatively, the solution showing the highest pay-off in the former population can be used, “**best**.” Other settings, including the previous solution of the target slot, have also been proposed. The second configuration parameter, y , denotes the number of difference vectors. While the original implementation—demonstrated above—used only a single difference “1,” it was later shown that by employing the sum of two difference ($y : “2”$) vectors, a stabler search direction can be achieved. The last parameter, z , specifies the crossover operator, for which there are two common variants: The uniform implementation as shown in the former algorithm, denoted “**bin**” (as it involves independent binomial experiments), and “**exp**,” a scheme based on an exponential random experiment, which is comparable to the two-point crossover of genetic algorithms, explained later. In the “exponential” scheme, a crossover point is uniformly randomly selected. The length is determined by a Bernoulli process starting from the crossover point and terminating at the position where the Bernoulli trial fails. The expected value for the trial is again the user-defined crossover rate (CR). The solution vector is considered cyclic in

this process so that if the crossover section exceeds the vector length, it continues at the beginning of the vector. One explanation for why this scheme is called “exponential” can be given when examining the probability of the crossover length l , which has the following probability of being equal to n ,⁷

$$P(l = n) = \begin{cases} \text{CR}^{n-1}(1 - \text{CR}) & n \in [1, N] \\ \text{CR}^{n-1} & n = N, \end{cases}$$

where N is the size of the population and, as before, $\text{CR} \in [0, 1)$ is the crossover rate. For the expected length, we have that

$$E(l) = \frac{1 - \text{CR}^N}{1 - \text{CR}}.$$

The success of differential evolution strongly depends on a proper choice of its control parameters, the mutation strength (F) and the crossover rate (CR) (Liu, 2002; Gämperle et al., 2002). Similarly to other evolutionary algorithms, numerous alternative schemes for mutation, recombination and selection, and also adaptive and self-adaptive variants have been proposed. Additionally, different approaches for alternative problem formulations, for example, incorporating constraints and multi-objective methods, have been developed (for a detailed discussion on related topics, see, for example, the contributions in Chakraborty, 2008).

4.1.4 Genetic Programming

Similarly to evolutionary programming, genetic programming (GP) is a technique to evolve programs rather than optimizing functions (Koza, 1992, 1994; Jacob, 2001; Poli et al., 2008). Its main application is hence for machine learning problems, in which it can, for example, be used for the evolution of hypotheses or programs. Consequently, the search space is given by objects that reflect the flow of these programs. In its classical form, GP uses trees, exploiting the fact that any program can be principally so represented. GP uses the same class of operations found in other evolutionary algorithms, including mutation, recombination and selection, all of which have to be designed to operate on trees. This is schematically demonstrated in Figure 4.2; the programs in this example are terms consisting of binary arithmetic operators. In general, any operator, function, relation or control statement with any input arity (but commonly limited to a one-dimensional output) can be applied, leading to an arbitrary tree morphology. In the GP terminology, all expressions taking arguments are called functions. The set of terminals, that is, expressions that do not take arguments and that are hence not expandable, can include numbers, variables or words. GP expressions hence consists of functions whose arguments are terminals or again functions. In order for genetic programming to evolve valid programs or structures, the used symbols must meet two conditions (Koza, 1992, pp. 81): (1) *Type consistency*: Input parameter types and function return types must correspond for all permitted combinations and contexts. (2) *Evaluation safety*: The evaluation of all terms must yield a well-defined state. Especially, no uncaught exception, errors or infinite loops may occur. These conditions, termed *closure* can be ensured implicitly (for example, using implicit type conversions or penalties for terms yielding errors) or explicitly (e.g. by restricting the symbol set to only yield consistent expressions or by the use of strict typing). Another requirement is that the symbol set must be sufficiently “complex” so that there exists a combination of these symbols that is a solution to the approached problem (Koza, 1992, pp. 86). The crossover operation illustrated in Figure 4.2, termed subtree crossover,⁸ mixes two expressions in the following way: In each parent solution, a crossover point, a node in the tree, is randomly selected. The subtrees defined by the selected crossover nodes are then exchanged, resulting in two offspring individuals.

A number of alternative crossover operators have been suggested (for an overview, see Poli et al., 2008, pp. 44–46). In the *one-point crossover* (Poli and Langdon, 1997), crossover points are selected

⁷For details, see Zaharie (2007) and Dasgupta et al. (2009).

⁸Here, we present a variant yielding two offspring solutions. The more common one-child version is obtained by simply omitting the mixing steps involved in the second child.




Figure 4.2: Genetic programming crossover operation. Parent individuals (a) and (b) are recombined to form two offspring solutions (c) and (d). As crossover points, nodes are randomly selected in the parent individuals. The subtrees spawned by these nodes (highlighted in gray) are then interchanged. In this simplistic example, the functions are the binary operations “+”, “−” and “ \cdot ” (written as “ \times ”, for better readability), with arguments denoted as capital letters. The arithmetic expressions the tree representation correspond to are shown in polish notation (which, for example, corresponds to functional programming language as LISP) above and below, respectively.

such that the donated subtree of one parent is located at a similar position as the subtree to be replaced in the other parent. The rationale behind this and other related operators is the reduction of bloat, the non-productive growth of trees. A higher degree of mixing is achieved by employing *uniform crossover* (Poli and Langdon, 1998), which, similarly to the correspondent operator in genetic algorithms (see Section 4.6.2), traverses the tree and randomly exchanges subtrees with a uniformly distributed probability.

In his original work, Koza (1992) dispensed with the use of any mutation operator. This is opposed to EP, where initially mutation was the only genetic operator. Nevertheless, today a multitude of different mutation operators have been proposed and are commonly used in practice (Poli et al., 2008, pp. 42–44). A straightforward approach is to randomly select a subtree in the parent individual and replace it with a different, randomly generated subtree. In the “*headless chicken*” operator, an individual is randomly created, and the mutation operation is implemented as a crossover between the parent and this random candidate (Angeline, 1997). Point mutation is performed in a similar way as in genetic algorithms (see Section 4.7): The parent’s tree is traversed, and with a pre-defined probability, each node is replaced with a compatible random symbol.

In contrast to other evolutionary approaches, where often crossover and mutation are executed consecutively, GP normally uses only one operator—either mutation or recombination—at a time. Traditionally GP makes use of the same selection schemes as GAs (Section 4.5).

The fitness function is strongly problem-specific. It is often implemented as a test that executes the candidate program and infers and evaluates its results.

In comparison with other evolutionary approaches, the random initialization of candidate solutions is less straightforward, since it involves the construction of trees with potentially unlimited depths. It is common practice, however, to limit the depth of initial trees to a predefined size and randomly construct full or incomplete trees with or up to this size. Two alternatives, specifically aimed at reducing bloat, have been proposed: *uniform initialization* (Langdon, 2000) and *Lagrange initialization* (Dignum and Poli, 2007).

The nonlinearity and the unbound length of the tree or graph representations used by GP allows for a considerably higher flexibility compared to other evolutionary algorithms. On the other hand, this property also poses a severe problem in that it prohibits unique solutions and may lead to computationally expensive structures. This has been partly addressed by an alternative linear chromosome representation devised by a number of authors (e.g., see Banzhaf, 1993; Nordin, 1994). The authors argue that the so-modified representation does not actually impose any restriction when applying GP to the automatic generation of programs, which are predominantly stored in an automaton in a sequential fashion.

Although long identified, the cause of non-functional growth of GP chromosomes, called bloat, present especially when using tree representations, is still a subject of intensive research (Poli et al., 2008, pp. 101–104). Different techniques have been proposed to reduce bloat (Poli et al., 2008, pp. 104–108), including the incorporation of size-reducing operators such as deletion, constraining the depth of trees, for example, by adding a fitness bias toward shorter programs or by employing tailored operators preventing bloat.⁹

Another common approach to reduce the time and space complexity of evolvable programs is to encapsulate common subtrees. Angeline and Pollack (1992) and similarly Rosca and Ballard (1994) devised a routine to automatically detect and store such subtrees in a common data base. Specific mutation operators were provided to insert a corresponding subroutine call into a parent’s program. A similar approach, called *automatically defined functions* was proposed by Koza (1994), albeit storing the subtrees with the candidate solutions instead of in a central data base.

Comparably to other artificial evolution approaches, GP depends on a number of control parameters, including population sizes and operator probabilities. A number of adaptive and self-adaptive strategies have been proposed (Angeline, 1996; Teller, 1996; Edmonds, 2001, for example,). In addition, the

⁹A recent rationale for the latter is the conjecture that genetic operators, especially crossover, lead to a skewed tree depth distribution of exceedingly short candidates with below-average pay-offs and significantly longer programs with above-average fitnesses (Dignum and Poli, 2007). Following this argument, bloat can be characterized as a GP-specific form of genetic drift.

same control parameter adaptation strategies as in other evolutionary algorithms can be employed; see Section 4.8 and Appendix E for details.

In addition to computational intelligence applications, GP has been successfully applied to a number of structural optimization problems as found in design tasks, for example, for integrated circuits.¹⁰ A comprehensive, yet of course non-exhaustive, overview was compiled by (Poli et al., 2008, pp. 111–130). A noteworthy specific class of applications is the developmental genetic programming approach in which programs are evolved that themselves produce the actual goal structures. GPs can, for instance, be used to identify rule sets for cellular automata (Andre et al., 1996) or evolve L-systems (Lindenmayer, 1968a,b) simulating the growth of plants (Jacob, 2001, pp. 449).

4.1.5 Artificial Immune System

An *artificial immune system* is an approach aimed at learning and memorizing responses to entities or properties that it is exposed to, typically employing simple agents (Dasgupta and Niño, 2008). Although a number of variants have been proposed as early as in the late 1980ies, there in the field of machine learning (e.g., Farmer et al., 1986; Bersini and Varela, 1991), it became popular also for optimization applications with the introduction of an algorithm based on the clonal selection operator, proposed by De Castro and Von Zuben (2000) and later, using a slightly different approach, by Cutello and Nicosia (2002). This variant makes use of two populations: an antigen population and the antibody population, which is evolved to exhibit an adequate response to the antigens. The fitness of an antibody candidate is determined by its antigenic affinity, that is, its ability to respond to the antigens. In each iteration, a number of clones of antibodies are generated. Each antibody is replicated with a proportion reflecting its fitness. The so selected solutions undergo hypermutation, with a mutation rate reciprocal to the affinity a candidate exhibits. When applied to optimization, the antigen population is typically replaced with the function to be optimized. Affinity can then be defined as the function value resulted from the decision variables represented by the antibody.

4.1.6 Swarm Intelligence

*Swarm intelligence*¹¹ summarizes approaches that mimic the social behavior observed in species such as fish, birds or insects. In the context of computation, it typically describes systems employing populations of simplistic entities, called micro-agents or sometimes “boids” (Reynolds, 1987). Each agent follows a small set of simple rules. The collective, however, through interactions of their agents with the environment and with each other, can be employed to solve large-scale, for example, optimization problems. Although based on a different role model or analogy, swarm intelligence optimization approaches have a number of concepts in common with evolutionary algorithms. Recently, a vast number of optimizers following the swarm intelligence paradigm have been proposed, a selection of which are briefly summarized in the following:

One of the successful attempts to employ swarm intelligence for optimization is the particle swarm optimizer (PSO) approach (Kennedy and Eberhart, 1995). Similarly to evolution strategies and differential evolution, the decision space consists of n -dimensional vectors, where each component denotes one decision variable. The agents or candidate solutions are often referred to as particles, “living” in the corresponding hyperspace; the pay-off of a particle’s position is determined by an objective function. A flock-like behavior is modeled by moving particles toward

- (1) the globally best paying off solution observed so far and
- (2) the best configuration the specific particle has visited.

In order to increase the abidance probability for unexplored search space regions, both components of the direction vector are scaled with a random variable, in other words, incorporating a fuzziness when determining the new position. The overall update process of one candidate can be described as a two-step routine, in which first the direction vector, often termed velocity vector, and then the position

¹⁰For example, through the evolution of domain-specific language programs.

¹¹Commonly, Beni and Wang (1989) are credited for the introduction of this term into the field of robotics and artificial intelligence.

of the particle are computed. Formally, given $\mathbf{v}_i^{(k)}$ the current velocity vector of a particle with decision vector \mathbf{x}_i , a memorized best solution \mathbf{x}_i^* of this particle and a globally best observed solution, \mathbf{g}^* , the new velocity $\mathbf{v}_i^{(k+1)}$ is defined as:

$$\mathbf{v}_i^{(k+1)} := \mathbf{v}_i^{(k)} + \alpha \mathbf{u}_1 \circ (\mathbf{x}_i^* - \mathbf{x}_i^{(k)}) + \beta \mathbf{u}_2 \circ (\mathbf{g}^* - \mathbf{x}_i^{(k)}),$$

where $\mathbf{u}_j (j = 1, 2)$ are two independent uniformly distributed random vectors¹², and α and β are user-provided parameters, determining whether the new position is directed more toward the globally best solution or the memorized local one. These parameters are strongly problem-specific and thus also crucial for the convergence behavior; in analogy with tabu search, the local contribution can be broadly associated with an intensified exploitation of the candidate's neighborhood, while the global one will lead to a higher diversification. Using the previously evaluated direction $\mathbf{v}_i^{(k+1)}$, we then have for the actual position update step that,

$$\mathbf{x}_i^{(k+1)} := \mathbf{x}_i^{(k)} + \mathbf{v}_i^{(k+1)}.$$

PSO and its numerous variants have been successfully applied to a wide range of problems (see, for example, Sedighizadeh and Masehian, 2009; Eberhart and Shi, 2001).

Other swarm intelligence optimization approaches include *ant colony optimization* (Dorigo, 1992): Ants are known to “communicate” through pheromones, for example, triggering a mass recruitment for the exploration of foraging areas. Pheromone concentrations accrue with the amount of ants excreting pheromones at the same locations. In addition, the pheromone evaporates over time, leading to a decrease of its concentration. The combination of both effects leads to an increasing preference of ants for shorter paths. Ant colony optimization makes use of this fact, rendering it particularly suited for routing and scheduling problems.

More recently, the fast-emerging field of swarm intelligence-based approaches has brought forward a number of other remarkable concepts: The *bees algorithm* (Pham et al., 2006) and the *artificial bees colony* approach (Karaboga and Basturk, 2007) mimic the foraging behavior of bees, including their recruitment strategy based on dances. In the *firefly algorithm*, candidate solutions are attracted by light emitted from companion candidates (Yang, 2009). The brightness of a solution is proportional to its fitness. In order to propel the exploitation of the current neighborhood, the brightness of individuals exponentially decreases with distance. This concept serves a similar purpose as the velocity component given by the individual's memory (\mathbf{x}_i^*) in PSO and the intensification and diversification techniques in tabu search. A number of successful applications of the firefly algorithm, especially on highly multimodal and noisy search functions, have been reported (Yang, 2009, 2010; Chai-ead et al., 2011). In the *cuckoo search* approach, the parasitic brood behavior is modeled by allowing candidate solutions, cuckoo eggs, to be placed in different host nests (Yang and Deb, 2009). The quality of an egg—the fitness of the solution—determines its survival probability. Key to this straightforward strategy, which has proved to even outperform PSOs and GAs under certain conditions, is its use of Lévy flights (Viswanathan et al., 1999) for the random walks that determine the cuckoo movement when targeting a new nest.

4.2 Genetic Algorithm Flow

The flow of genetic algorithms does not significantly differ from other evolutionary algorithms and can be broadly characterized as an iterative application of the operators *selection*, *recombination* and *mutation*, followed by the evaluation of the pay-off of the newly generated offspring solutions (Holland, 1975; Goldberg, 1989). As schematically shown in Figure 4.3, these steps are repeated until a predefined termination criterion—often a maximum number of iterations (generations) or a stagnation of the best individual's fitness—is met.

Genetic algorithms typically use a fixed-size population. The initial population consists of randomly generated individuals. In addition, if already known solutions are to be taken into account, obtained, for

¹²As commonly scalars are used rather than vectors, in which case this parameter can be thought of as a random step size. Using the more general random vector implementation and the Hadamard product, also randomizes the direction and can hence be compared with mutation in evolutionary algorithms.

example, from a previous pre-optimization procedure, they can be included into the initial population. This insertion of pre-existing candidate solutions is often referred to as seeding. The initial solutions are evaluated using a problem-dependent objective function, whose result is translated into a pay-off value that characterizes the fitness of a solution. Pairs of solutions are randomly selected from the population. The probability for a specific individual to get selected is proportionate to the pay-off that it exhibits. In other words, fitter individuals receive a higher selection probability than those solutions that have a lower merit. The selected solution pairs then mate, forming two offspring individuals, which are randomly modified, with a low probability, through mutation. The so newly created solutions are evaluated thereafter, starting a new iteration step.

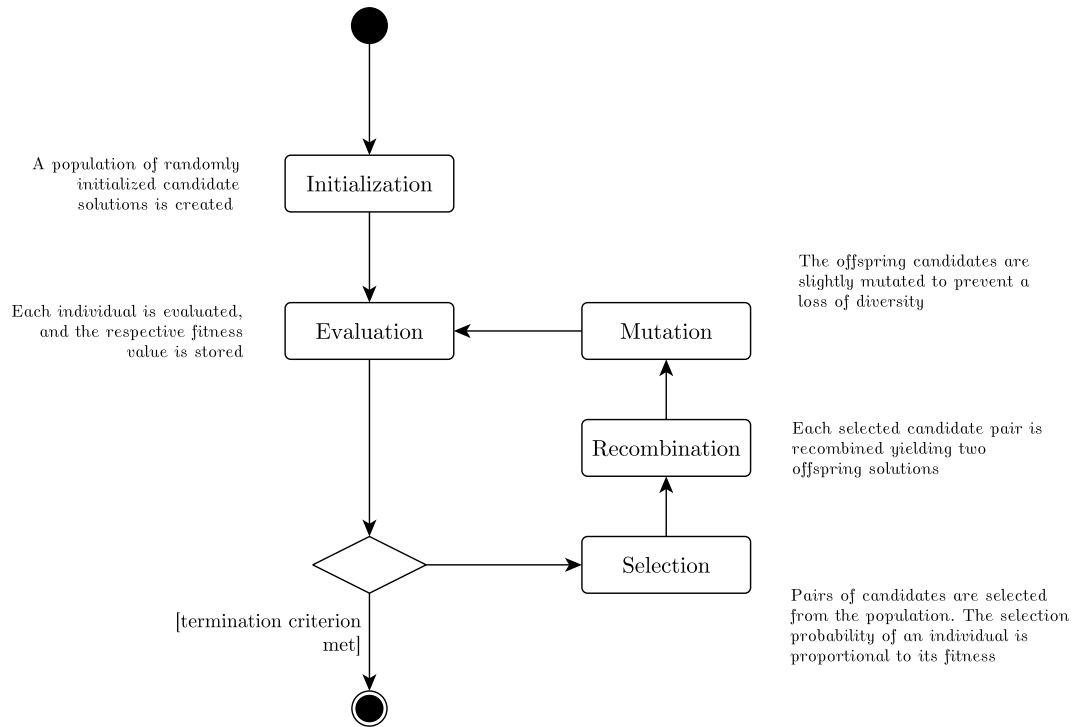


Figure 4.3: Operation flow of genetic algorithms.

4.3 Optimization Parameter Representation

More than all other artificial evolution approaches—with the exception of genetic programming—classical genetic algorithms employ a strict disjunction between the outward appearance (phenotype) of a candidate solution and its encoding (genotype).¹³ This is in accordance with the role model, the natural evolution process, and, more importantly, render GAs well suited for combinatorial problems. Conventionally, GA individuals consist of a chromosome, an array of genes whose alleles (values) are from a finite alphabet, an expression function translating the chromosome into a problem-specific representation and a pay-off function that potentially dynamically depends on the environment. The latter is often applied in order to suppress the occurrence of superindividuals—individuals whose pay-off outshines other candidate solutions at a very early stage—and is often referred to as scaling (Forrest, 1985).

¹³A detailed discussion can, for example, be found in the original work of Holland (1975) or in (Goldberg, 1989).

4.3.1 Binary Versus Real-coded Genetic Algorithms

Until the late 1980s, there was little common ground between the real-coded evolution strategies and the binary-valued genetic algorithms. By virtue of Holland's (1975) schema theorem (see Section 4.10.1 for a detailed discussion), the efficacy of GAs was attributed to the implicitly parallel processing of hyperplanes spawned by discrete alphabets of low cardinality. Moreover, GAs were not originally designed as function optimization but rather as simulation frameworks to study natural and artificial adaptive systems (De Jong, 1993). Evolution strategies—which were mainly advocated by European researchers¹⁴—on the other hand, were essentially considered to be sophisticated random walk approaches, in which information exchange through recombination plays only a marginal role. The emergence of real-coded genetic algorithm (RCGA) and their successful application to a number of real-world problems (for example, Davis and Coombs, 1987; Lucasius and Kateman, 1989; Janikow and Michalewicz, 1991; Michalewicz, 1996) drove the need to revisit the theory on information processing in GAs.

Goldberg (1991), for example, identified three problems when binary-coded genetic algorithm (BCGA) are applied to optimize functions with continuous domains:

- (1) An impaired hill-climbing behavior under mutation due to Hamming cliffs.¹⁵
- (2) Especially for non-multimodal search functions, the convergence efficiency may be inferior in comparison with real-coded GAs.
- (3) As will be discussed later, from the schema theorem, Goldberg derived a (widely-disputed) concept of what he (and others) term *deception* (see, for example, Goldberg, 1989, pp. 46–53), which renders certain classes of problems hard for GAs to solve. Goldberg (1991) claims that GAs that use alphabets with a low cardinality are more prone to deception than, for example, real-coded GAs.

Goldberg (1991) argues that a real-coded GA vastly concentrates its search on slices of the search space that exhibit a pay-off above average. These above-average slices then play a similar role as the hyperplanes spanned by binary-string GAs; they form a virtual alphabet from which the real-coded GA draws new solutions. Virtual characters hence play a similar role for real-coded GAs as schemata (the hyperplanes spanned by bit-combinations) for bit-string GAs. In contrast, however, the alphabet is chosen implicitly. Goldberg argues that this property appears advantageous at first sight but leads to side-effects. As an example, he shows an effect he refers to as blocking, in which an optimum is disregarded since it is not covered by previously constructed virtual characters.

An alternative framework for the analysis of real-coded genetic algorithms was developed by Radcliffe (1991). He proposes the use of equivalence classes that are determined by the effect the GA operators have on the chromosome. More specifically, he derives a number of criteria that the equivalence classes, termed *formae*, have to meet under selection, recombination and mutation in order to adequately describe the information propagation in GAs. Schemata then present a special case of these classes in the scope of binary alphabets.

Wright (1991) proposed a simple translation of the hyperplane processing of binary GAs into parameter intervals in real-coded GAs. Accordingly, he developed a real-coded GA variant in which binary mutation and recombination operations are emulated on the real-coded vectors. He showed that for a number of benchmark functions, a faster convergence and slightly improved final results could be achieved. The latter, can be attributed to the accuracy constraints given by the bit string representation. For the highly multimodal test function F_5 (Shekel's Foxholes; De Jong, 1975), however, the bit-string GA significantly outperformed the real-coded GA. Wright argued that the function is highly periodic and thus well-suited for a binary representation, albeit also for a 30-degree rotated version of F_5 the binary representation led to better results than the real-coded version. A similar interval schema theorem was later also proposed by Eshelman and Schaffer (1993). A comprehensive overview on the former and other analysis approaches—for example, based on the expected progress (Mühlenbein and

¹⁴With a number of exceptions, for example, Weinberg and Berkus (1971), as pointed out by Goldberg (1991).

¹⁵The standard binary coding does not accord with the continuous in terms of distance between neighboring values. For example, in a three-bit code, the numbers 3 and 4 do not share single common bit (011 100). Often an alternative bit-code, termed Gray code, which will be explained later, is applied to counteract problems associated with Hamming cliffs.

Schlierkamp-Voosen, 1993) or a discrete-time stochastic process (Qi and Palmieri, 1994a,b)—can be found in Herrera et al. (1998) and Deb (2001).

The success of real-coded GAs strongly depends on the selection of adequate genetic operators. In contrast to binary-coded GAs, which typically make use of only a small set of classical recombination and mutation operators (see Section 4.6 and Section 4.7), a multitude of crossover operators (for example, Janikow and Michalewicz, 1990; Michalewicz and Janikow, 1991; Wright, 1991; Radcliffe, 1991; Michalewicz, 1996; Eshelman and Schaffer, 1993; Mühlenbein and Schlierkamp-Voosen, 1993; Deb and Agrawal, 1994; Herrera, 1995; Voigt et al., 1995; Nomura and Miyoshi, 1996; Ono et al., 1999; Higuchi et al., 2000; Tsutsui and Goldberg, 2001; Deb and Anand, 2002; Ling and Leung, 2006; Deep and Thakur, 2007a; Li and Lu, 2007) and mutation operators (Janikow and Michalewicz, 1990; Davis, 1991; Michalewicz and Janikow, 1991; Mühlenbein and Schlierkamp-Voosen, 1993; Voigt and Anheyer, 1994; Mäkinen et al., 1999; Berry and Vamplew, 2004; Temby et al., 2005; Ling and Leung, 2006; Deep and Thakur, 2007b; Li and Lu, 2007; Korejo et al., 2010; Gong et al., 2010) have been proposed for real-coded GAs.

A comprehensive performance study comparing different real-coded GAs, with different crossover and mutation operators, and a simple binary GA were conducted by Herrera et al. (1998)—with later extensions (Herrera et al., 2003). It was shown that when applied as a function optimizer, a number of RCGA variants were able to outperform the simple BCGA both in terms of online and offline performance.¹⁶ However, the performance of RCGA strongly depend on the crossover and mutation operators employed. Although the studies are limited to function optimization problems whose variables are within a continuous domain, no distinct conclusion, leading to a one-best set of real-coded GA operator or an absolute advantage of real-coded GAs, can be drawn. In the regime of mixed optimization problems consisting of integer programming tasks or other combinatorial problems in conjunction with real-coded domains, an even more ambiguous picture can be expected.

As an alternative to applying real-coded genetic algorithms to pure function optimization problems, other authors have proposed to adapt the binary representation to this specific task. For example, Schraudolph and Belew (1992) argue that a number of efficiency issues of function optimizations using binary-coded GAs can be associated with the fact that at the beginning a GA is mainly concerned with finding optimal alleles for the high-significant genes while in a later phase a fine-tuning of the less significant positions is performed. He suggests that this may lead to *genetic hitchhiking*, as the bits of less significance may be arbitrarily linked to higher order genes. As a countermeasure—inspired by previous work by Shaefer (1987)—he proposes the use of a *zoom operator* acting on each parameter independently. The GA is started with a coarse chromosome resolution. As soon as a parameter is considered to have converged to values within a certain interval, all individuals are recoded such that the respective parameter is restricted to the narrower interval, leading to finer resolution there. Since the adaptation of the chromosome is influenced by phenotypic considerations, this approach can be considered a Lamarckian/Darwinian hybrid method. A different technique, also aiming at increasing the chromosome resolution in later GA stages was proposed by Whitley et al. (1991), called *delta-coding*: A conventional GA is run until a predefined diversity threshold is undermatched. The best solution is stored, and the GA is restarted; albeit, not with the original coding scheme but with a bit-string that represents the difference from the previously stored best solution. This procedure is repeated with increasingly finely grained deltas, in order to concentrate the search on sub-regions of the search space.

4.3.2 Other Representations

Other discrete representations using alphabets with a higher cardinality than two have also been proposed. According to Holland (1975), binary alphabets constitute an ideal representation for genetic algorithms as they facilitate their implicit parallelism. This statement has been challenged, for example, by Antonisse (1989), who offered an alternative interpretation of the schema theorem, contrariwise

¹⁶The terms *online* and *offline* performances have been first proposed by De Jong (1975). The *online performance* denotes the fitness average of all individuals after a given generation and is hence an assessment estimator of the entire population dynamics. The *offline performance* is the average of the best fitness values in the population up to a given iteration step. It is thus better suited to measure convergence.

suggesting that high-cardinal alphabets lead to an improved parallel information processing. Later Fogel and Ghoseil (1997), extending the findings of Battle and Vose (1993), presented arguments that invalidate the claims of advocates of both low and high cardinality alphabets, in basically showing that an ideal representation is both critical but also strongly problem-dependent. Users of non-binary discrete representations include Bramlette (1991), who applies an integer representation to optimizing maximum flow problem with a GA. Similar concepts were introduced for flow and job shop problems (see, for example Davis, 1985; Grefenstette and Gopal, 1985; Goldberg and Lingle Jr., 1985; Oliver et al., 1987; Mühlenbein et al., 1988; Fogel, 1988; Whitley et al., 1989; Banzhaf, 1990; Fox and McMahon, 1991; Nakano and Yamada, 1991; Seniow, 1991; Syswerda, 1991; Homaifar et al., 1993; Reeves, 1995; Su et al., 2009).¹⁷ Integer-valued vector and matrix representations were also proposed to solve linear and nonlinear transportation problems (Vignaux and Michalewicz, 1991; Michalewicz et al., 1991). Ichikawa and Ishii (1993) used a GA with an integer-string representation in training algorithms for artificial neural networks.

The multitude of different representations and the contradicting statements of their performance in GAs, gained both from empirical and theoretical considerations, greatly complicates an adequate a priori choice when confronted with novel optimization problems. This practical difficulty motivated Schnier and Yao (2000) to develop a hybrid approach in which individuals are encoded in multiple ways from which the GA automatically selects an appropriate representation (also compare Rothlauf, 2006). In their empirical investigations, they showed that such a technique may be suited to overcome performance issues that they associated with ill-selected representations.

4.3.3 Schemata

A common concept when illustrating and analyzing the processing behavior of (binary-coded) genetic algorithms is to consider the hyperplanes that partition the search space. In the binary representation case, the hyperplanes, called schemata, are given through the instances of equivalence classes that are defined by setting a number of positions to specific values and letting the rest of the chromosome take arbitrary alleles (Holland, 1975). The undefined positions are often denoted with an asterisk (*). For example, given a chromosome of length five, the schema $10*1*$ is defined as follows:

$$10*1* := \{10010, 10011, 10110, 10111\}$$

A straightforward asset of schemata is their fitness, for instance, in order to identify above-average partitions of the search space. Given a schema H , the schema fitness is obtained by evaluating the average fitness of the candidate solutions that it represents:

$$\hat{\mu}(H) = \sum_{A \in H} \mu(A) / |H|.$$

Figure 4.4 illustrates schemata for the maximization of the function:

$$\begin{aligned} \mu : D \times D &\rightarrow \mathbb{R}, \quad \text{where } D := \{d \in \mathbb{N} \mid 0 \leq d \leq 6\} \\ \mu : (x, y) &\mapsto (3 - |x - 3|)(3 - |y - 3|), \end{aligned}$$

which has a maximum value of 9, for $x = 3$ and $y = 3$.

Let \mathfrak{S} be the binary representation of the search space such that for each $A \in \mathfrak{S}$:

$$A := ((x)_b, (y)_b) := ((a_5 a_4 a_3)_b (a_2 a_1 a_0)_b),$$

where the subscript b denotes the standard binary representation and the genes a_i hence take the alleles 0 or 1. Since $(x, y) \in [0, 6] \times [0, 6]$,

$$\mathfrak{S} = \{A = (a_5 a_4 \dots a_0)_b \mid a_i \in \{0, 1\} \wedge \neg((a_5 = a_4 = a_3 = 1) \vee (a_2 = a_1 = a_0 = 1))\}$$

¹⁷A comprehensive review on the different approaches was given by Larranaga et al. (1999).




Figure 4.4: Example of Schemata.

The average fitness of the search space is,

$$\hat{\mu}(\mathfrak{S}) = 2.25.$$

In Figure 4.4, two schemata $H_1 : 1*0|***$ and $H_2 : ***|00*$ are demonstrated. The former (marked with blue circles) traverses the entire y axis in two x values (4 and 6), while the latter (marked with red triangles) takes arbitrary x values for two fixed y values (1 and 2). The fitness of H_1 is about 1.2857; the fitness of H_2 is about 0.6429. Neither schema is above average, and they can be expected to have a low survival probability. The same is true for the combined schema,

$$H_{12} = 1*0|00* \equiv \{(4, 0), (4, 1), (6, 0), (6, 1)\},$$

which averages to 0.5.

An example for an above-average schema is $***|01*$ with a fitness of about 3.21. It also happens to be one of the 2^5 schemata whose instance is the maximizer $011|011$. Generally, each candidate of a problem with a chromosome length of l is an instance of the $2^{(l-1)}$ schemata whose defining positions coincide with the corresponding alleles of that candidate. Broadly speaking, genetic algorithms can hence be considered as an implicitly parallel search and recombination process of above-average schemata. We shall see later (Section 4.10.1) that this conjecture constitutes the basis for classical GA theory.

4.3.4 Gray Code

As mentioned before, one of the problems associated with standard binary-coded GAs applied to function optimization is its non-continuous neighborhood relation often referred to as Hamming cliffs or walls: Values adjacent when represented as decimal numbers exhibit a large bit difference in the binary domain. For example, consider the simplistic one-dimensional maximization problem shown in Figure 4.5(a), where the maximum is at a parameter value of 4, from which one of its adjacent values, 3, exhibits the maximum Hamming distance (bit difference) of 3: applying, for example, point mutation to this solution candidate will most likely result in a divergence from the optimum. In order to alleviate this problem, the use of a reflected binary code has been proposed (Caruana and Schaffer, 1988). In Gray code, originally developed to reduce errors for analog-digital interfaces (Gray, 1953), numbers

adjacent in the decimal system have a bit difference of one in the binary representation. The Gray code for n bits can be constructed using the following recursive formula:

$$\begin{aligned} b^{(0)} &:= () \\ b^{(n+1)} &:= (0 \boxplus b_i^{(n)}, 1 \boxplus b_i^{(n)}) \quad i = 1, \dots, 2^n, \end{aligned}$$

where $b^{(n)}$ is the sequence of bit strings with length n , and \boxplus denotes the concatenation operator, e.g.,

$$1 \boxplus 010 = 1010.$$

As shown in Figure 4.5(b), in contrast to the standard binary encoding, the Hamming distance between the Gray-represented solution candidates is equivalent to the decimal distance.

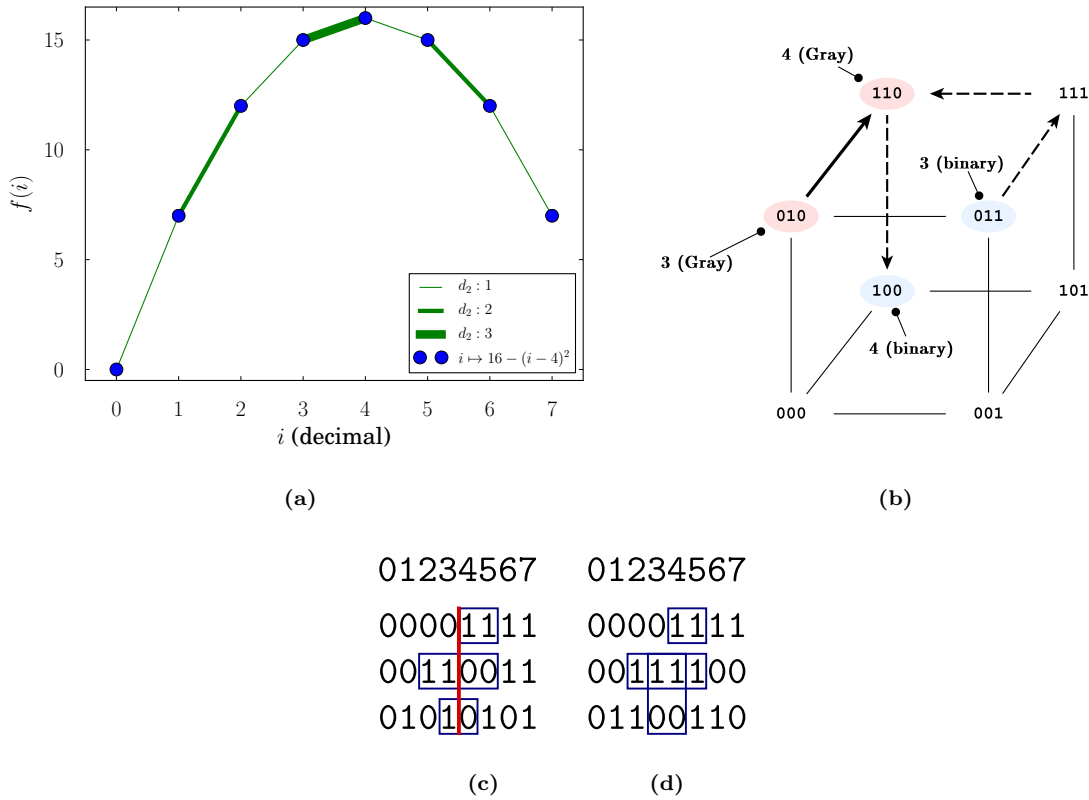


Figure 4.5: Comparison of binary and reflected Gray code. (a) Function graph where the width of the edges illustrate the Hamming distance (d_2) between the points in the standard binary domain. (b) Standard binary and Gray code representations of the numbers 3 and 4: while in the binary representation the path between both numbers involves three nodes, the Gray code representations are adjacent. (c) and (d): Comparison of schemata representing the optimum solution in the binary and the Gray code domain; The wide red line (c) illustrates the Hamming cliff of the standard binary representation.

The resulting advantage for function optimization can also be illustrated when comparing the schema mixing properties of the standard binary (Figure 4.5(c)) and the Gray-coded (Figure 4.5(d)) representation cases.

- The average fitness over all solution candidates is 10.5.
- For both representations, the maximizer is an instance of the schema 1^{**} , which has the same fitness in both cases.

- The schema with the second allele of the maximizer being the only defining position differs: It is $*0*_b$ in the binary code case and $*1*_g$ in the Gray code case. Also their performance differs. In the standard binary case, the schema is below average,

$$\hat{\mu}(*0*_b) = 9.5,$$

while the respective Gray-coded schema has an above-average fitness,

$$\hat{\mu}(*1*_g) = 14.5.$$

- The low bit alone contains little information and is below average for both representations. The second-order schema, defined by the second and the last position, again reveals the advantage of Gray code over standard binary representation. In the binary case, the maximizer is represented by $*00*_b$, which has a fitness of 8, in the Gray-code case, the respective schema is $*10*_g$, having a fitness of 15.5.

Although simplistic and incomplete, the former example demonstrates the potential benefit of Gray over standard binary code in the case of function optimization problems. By virtue of the NFL theorem (Wolpert and Macready, 1997), any general claim that Gray code outperforms standard binary representations must be considered false. Nevertheless, Whitley (1999) showed that for function optimization problems involving search spaces with bounded complexity, Gray code does indeed lead to a reduced number of optima and can hence be expected to improve the convergence behavior.

4.3.5 Polyploidy

Commonly, genetic algorithms utilize single-stranded chromosomes. Following the natural role-model, researchers have also proposed the use of di- or polyploid representations. That is, the genotype consists not only of one but of two or more chromosome strands. Each gene may then take multiple alleles. In order to obtain a phenotypic representation of the candidate solution, a mapping, termed expression, determining the valid allele has to be applied. A common approach is to employ dominance. Similarly to Mendel's genetic model, an allele is said to be dominant, if it is the result of the expression in the case of heterozygosity; it is recessive if it results only if it is homozygote. The main application of di- or polyploidy in artificial evolution approaches is to increase the adaptability of solutions to time-varying conditions. Dominant alleles are expressed with a higher preference; recessive alleles are not lost, however. If the conditions change such that recessive gene values result in a higher pay-off, these alleles may become re-activated. Commonly, this approach is supplemented by a co-evolution of the dominance function that determines if an allele is dominant or recessive. An early approach—restricted to stationary problems—was introduced by Bagley (1967). There, the diploid chromosome is extended with a third, co-evolvable strand, representing the dominance value for each locus. Similarly, Hollstien (1971) proposed to use a dominance parameter for each locus. The resulting behavior can be simplified by the introduction of ternary alphabets, consisting of the three letters 0, 1_0 and 1.¹⁸ Both 1_0 and 1 are expressed as 1. Though, while 1 dominates 0, 1_0 is dominated by 0. Thus a change in dominance can be achieved by replacing 1_0 with 1 and vice versa.

Another study on dominance was performed by Brindle (1981). She explored different schemes to determine dominance values, including a randomly initialized dominance map, which is fixed throughout a GA run, schemes accounting for the current allele distribution in the population, and a scheme that makes use of a third chromosome that represents the evolvable dominance map. This last chromosome is treated independently of the diploid chromosome.

The restriction of these early attempts to stationary problems was criticized by Goldberg and Smith (1987), who argued that the application of polyploidy and dominance is best suited for time-varying problems and should hence be concentrated on those. This claim was empirically verified for time-varying knapsack problems, where a diploid chromosome and a tri-allelic dominance map (as proposed by Hollstien (1971)) resulted in a significant improvement compared to a haploid GA (also compare Smith and Goldberg, 1992).

¹⁸compare Holland (1975) and Goldberg (1989).

As an alternative to the previous dominance maps, Yoshida and Adachi (1994) developed a crossover operation between the homologous chromosomes, which they termed pseudo-meiosis and verified on a non-stationary traveling salesman problem (TSP). Greene (1994) propose to independently express and evaluate each homologous chromosome. The resulted fitness was then used as the dominance function. Collingwood (1995) applied a polyploid chromosome approach to a stationary transport scheduling problem.

An additive dominance scheme for diploid di-allelic genotypes was proposed by Ryan (1994). He introduces four alleles representing increasing integer numbers. A gene is expressed as one if and only if the sum of the alleles at the respective homologous loci exceeds a predefined threshold. Another approach using four alleles was developed by Ng and Wong (1995). Their allele set consists of a dominant and a recessive one, and the respective zero alleles. If a locus exhibits a dominant and a recessive gene the former is expressed. Otherwise the allele to is chosen at random. Lewis et al. (1998) show that a successful application of the former approaches to highly non-stationary search problems depends on a co-evolution of the dominance function, with which both approaches can be supplemented. Alternatively, the dominance map can be obtained inductively by relating the diploid genes to the phenotypic performance of the previous population. This procedure was proposed by Uyar and Harmanci (1999).

4.3.6 Genders

In contrast to nature where the paramount of sexual reproduction processes involves different genders, they play only an inferior role in GAs. Goldberg (1989), for example, considers the use of sexual differentiation in GAs as a specialization of niching (a diversity preservation technique we will introduce later). Ronald (1995) and Ratford et al. (1997) have introduced selection mechanisms based on seduction functions, mimicking the nature-observed selection of mating partners. In an approach by Goh et al. (2003), the population is artificially subdivided into male and female candidates. While all females are unconditionally selected for mating, males undergo a tournament selection procedure. Sánchez-Velazco (2003) introduced a gendered genetic algorithm that employs different fitness measures for male and female individuals. While males are selected competitively based on their objective value, the fitness of females is determined by age. A similar approach was put forward by Zhang et al. (2009). In addition to a mating restriction mechanism and a life-span fitness contribution, their algorithm is aimed at modeling the Baldwin effect.¹⁹ Recently, Zegordi et al. (2010) developed a gendered genetic algorithm to solve a distributed scheduling problem. Similarly to the previously described hybrid representation scheme proposed by Schnier and Yao (2000), their respective male and female chromosomes represent different aspects of the optimization task (for example, the former, the job assignment schedule and the latter a job priority table). During crossover, which involves a male and a female individual, first the target gender is determined and the parent with the opposite gender is converted to this target gender, to finally perform a simple uniform crossover. Another application of multi-gender genetic algorithms is for multi-objective optimization problems (see Section 4.11). Allenson (1992), Lis and Eiben (1996), and Esquivel et al. (1999), for example, associate each individual objective value with one gender.

4.4 Constraint Handling

Constraints are an important and critical aspect in optimization. In many cases the constraint satisfaction poses a profoundly numerically complexer problem than the optimization of the objective itself. The integration of constraint handling mechanisms has hence been a topic of evolutionary algorithm research for several decades. It also plays an important role in the optimization problems presented in this work, for example, in order to ensure the manufacturability of photomasks.

Constraints can be classified in a various ways, for example (Nocedal and Wright, 1999): (1) By their numerical complexity, as linear or nonlinear, (2) as equality and inequality constraints, and (3) by their overall computational complexity, as implicit or explicit constraints. The latter is a rather weak

¹⁹Broadly speaking, the Baldwin effect can be characterized as a selection criterion arising from phenotypic adaptation to environmental circumstances.

criterion. In real-world applications one could, for example, distinguish between constraints that can be evaluated in a straightforward, e.g., analytical way, and those that require time-expensive numerical solutions.

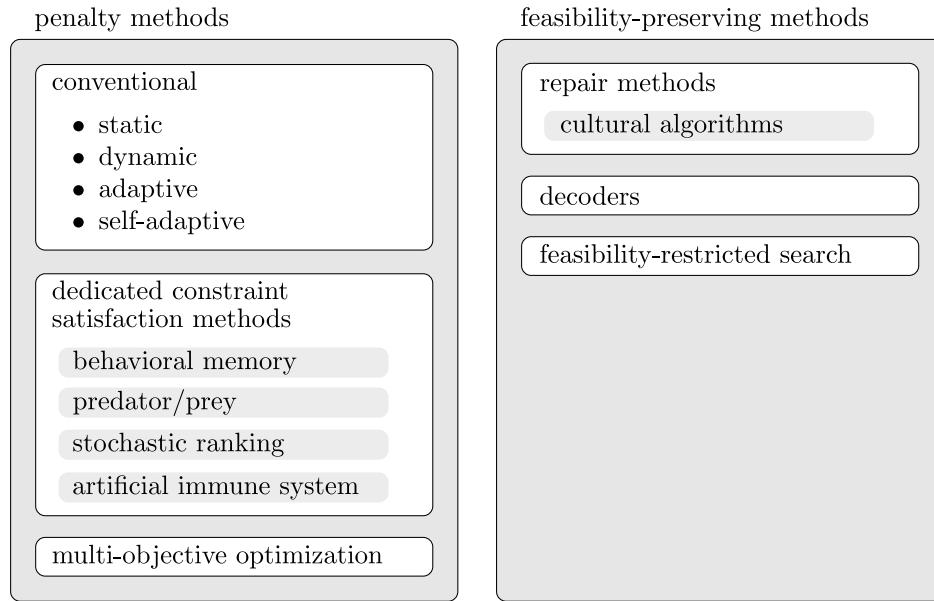


Figure 4.6: Taxonomy of constraint handling approaches in evolutionary algorithm based on Michalewicz and Schoenauer (1996) and Coello Coello (2002).

Because of their use of finite alphabets, standard genetic algorithms can be characterized as box-constrained optimization approaches. In general, however, evolutionary algorithms do not inherently ensure the feasibility of solutions but require additional techniques to be incorporated. Figure 4.6 provides a taxonomy of the state-of-the-art.²⁰ In general there are two methods by which constraints can be employed into EAs: (1) Using *penalty methods* to employ a preference of feasible over infeasible solutions or (2) by *feasibility-preserving measures* that ensure that at any time all solutions are feasible.

In the following, we will briefly summarize the most important techniques. Detailed surveys and discussions can be found in Michalewicz and Schoenauer (1996) or Coello Coello (2002).

4.4.1 Penalty Methods

Penalty methods are a common technique to solve constrained problems with unconstrained optimization algorithms. Assuming a minimization problem, a straightforward approach is to add the penalty values in form of a weighted sum to the objective value. A penalty value reflects the degree by which a constraint has been violated, while the weights are used both to normalize the penalties and to balance penalties and the objective. Penalty methods can be applied either as an *exterior point* method or as an *interior point* or *barrier* method. The former is primarily aimed at guiding infeasible solutions toward feasible regions; The latter, the barrier method, assumes feasible start points, for which the penalty function ensures their not breaching the feasibility boundary. Since the search for feasible points may be a numerically complex task in itself, it is often not eligible for population-based search approaches. An appropriate weight configuration is critical (Davis and Coombs, 1987; Richardson et al., 1989; Smith and Tate, 1993): Too high an uprating of penalties will restrict the search to the constraint satisfaction problem, whereas too low a value may cause the final solutions to be infeasible. To alleviate this

²⁰This taxonomy is based upon but slightly different from the ones presented by Michalewicz and Schoenauer (1996) and Coello Coello (2002).

problem, a number of dynamic, adaptive and self-adaptive weight control methods have been proposed, typically at the cost of (potentially significant) higher computation times.

As an alternative, a number of techniques have been proposed that separate the constraint satisfaction problem from the actual optimization. For example:

- *Behavioral memory*: In a sequential process, an increasing number of constraints are solved for, one at a time. Only in a final stage, the objective value is taken into account and optimized for.
- *Predator/prey*: The evaluation of the degree to which a solution is feasible and the complexity of each constraint is dynamically determined through co-evolution.
- *Stochastic ranking*: The fitness of individuals is determined either by its objective value or the degree of constraint violation. The switching between the two modes is performed probabilistically.
- *Artificial immune system*: Representative solutions known to be feasible are selected. Their constraint violation components serve as antigens, to which the non-feasible solutions are adapted by means of hypermutation and clonal selection.

A third approach is the use of multi-objective techniques (see Section 4.11), in which the problem is posed as compromise between the actual optimization and the constraint satisfaction problem. Numerous variants of this technique exist, as shown in Coello Coello (2002).

4.4.2 Feasibility-preserving Methods

In contrast to penalty methods in which constraint-violating candidates are principally permitted, *feasibility-preserving methods* aim at maintaining only feasible solutions. In *repair methods* this is achieved by transforming infeasible solutions into feasible ones. When employing *decoders*, the search is performed on an intermediate decision space. Points in the actual search space are obtained by a mapping from this intermediate space into the actual search space. This mapping (decoding) is to be defined such that it yields only feasible candidates. If the feasibility regions can be determined a priori, a *feasibility-restricted search* can be performed, for example, by employing an according genotypic representation or by using closed genetic operators, that is, operators that generate only feasible candidates. Another method to prevent operators from yielding infeasible solutions is to supplement them with a local constraint satisfaction search. This method can be applied even if feasibility regions are not obtainable in straightforward manner.

4.4.3 Selection of an Appropriate Method

The selection of an appropriate constraint handling technique is strongly problem-specific. If, for example, constraints can be determined in advance, feasibility restrictions display an efficient method in that they compress the search space to only the relevant partitions. Unfortunately, many problems are not amenable to an analytic approach, such that one has to resort to numerical alternatives. Conventional penalty methods, for example, often constitute an appropriate choice since they can be considered proven and robust, with many successful applications to a large variety of different problems. Their main drawback is their dependence on an empirical determination of weight factors. Because of the dramatically increased number of iterations, however, adaptive and especially self-adaptive penalty methods often do not offer a suitable alternative (Coello Coello, 2002). In many optimization problems, like the ones in this work, constraint violations can be distinguished between hard constraints, whose violation renders solutions invalid, and weak constraints, which, if violated, degrade but do not nullify the merit of solutions.²¹ Multi-objective approaches are especially well-suited to deal with these kinds of problems but have also been successfully applied to solve hard constrained problems. One of the reasons is that the problem formulation as a compromise between optimizing the objective and minimizing constraint violations typically causes the search to approach the constraint boundaries from both the feasible and infeasible side. This also leads to a great variety of final solutions, potentially ranging from ones that completely satisfy all constraints to others that do exhibit constraint violations but at the same time show a higher pay-off.

²¹The latter can be said to be a regularization framework, applied in order to reduce multimodality in the solution to an inverse problem.




Figure 4.7: This schematic of a search space S with two disjoint feasible regions \mathcal{F}_1 and \mathcal{F}_2 illustrates five different constraint handling methodologies: (1) *exterior point methods* such as penalty methods, in which candidates are guided toward feasible regions through penalty values, (2) *interior point methods*, in which feasible solutions are prohibited to leave feasible regions, (3) *repair methods*, which alter candidate solutions such that they become feasible, (4) *decoder methods*, which operator on a replacement search space from which a mapping into the feasible regions is performed, and (5) *feasibility-restricted search methods*, which commonly employ problem-specific knowledge in order for newly created solutions to be restricted to feasible regions or the feasibility border.

An additional issue has to be considered when the objectives and the constraint violation evaluations are results of a partially correct algorithm, which, for example, under certain inputs does not terminate.²² This can be considered a common situation, especially when various numerical models are involved as in the case of this work. Then, it might be necessary to precondition the inputs to the evaluation routines, employing, for instance, repair methods or decoders.

4.5 Selection

Briefly, selection can be characterized as the allocation of trials that undergo recombination and mutation, the result of which is evaluated and potentially replaces former candidates in the target population. In terms of function, the selection process is to choose individuals such that the product by means of genetic operations exhibits progressively increasing fitness. There, two contradicting properties can be distinguished: (1) In order to achieve a fitness progression, selection is to prefer solutions that exhibit an above-average fitness in comparison with other solutions in the population. (2) On the other hand, completely neglecting below-average solutions decreases the pool of genetic material, the number of alleles present in the population, and can hence lead to a loss of diversity. In other words, selection has to ensure that the expectation of trials for averagely fit individuals is one, significantly greater than one for the best individual, and not zero for individuals with a fitness below the average.

Traditional genetic algorithms use a non-overlapping *replacement* scheme. In each generation, the population is newly created from the individuals that resulted from the selection, recombination and mutation process. The generational gap is one. A typical extension to this scheme is to allow the best—or a number of highly fit—individual to be copied into the next generation without modification (De Jong, 1975). This technique is referred to as *elitist selection* and is applied to prevent the extinction of solutions with a superior fitness or their genetic material. With elitist selection the generational gap is $N - n_e/N$, when taking N as the population size and n_e as the number of elitist solutions.

Another deviation from a generational gap of one is given in *steady-state genetic algorithms* (Whitley and Kauth, 1988). In this variant, only the worst individual (or a fraction of the individuals with below-average fitness values) is replaced per generation. The turnover time per generation is hence significantly shorter than in standard GAs, normally resulting in a greater number of generations. Goldberg and Deb (1991), comparing different selection schemes, associate high growth rates of the best individual with this technique. Although this property may to a shorter convergence time than other approaches, it also bears the danger of a premature loss of diversity.

In contrast to evolution strategies (Section 4.1.1), in genetic algorithms, the offspring solutions resulting from recombination and mutation are directly transferred into the target population. Consequently, the number of individuals to be replaced in the population is equal or smaller than the number of offspring individuals. In case excessive individuals are generated or if the newly generated individuals are to compete with solutions from the old population, an additional selection step has to be performed, which is termed *survivor selection* to distinguish it from the *mating selection*. Both steps utilize identical (or at least very similar) selection schemes.

The most important selection schemes are briefly explained below.

4.5.1 Stochastic Sampling With Replacement

This technique is often termed *roulette wheel selection* since it can be illustrated by a roulette wheel or wheel of fortune that is subdivided into as many segments as there are individuals in the population. The selected individual is determined by spinning the wheel and choosing the solution corresponding to the segment at which the wheel stops (see Figure 4.8(a)). In order for the selection probability to reflect the fitness of an individual, the sizes of the segments have to be defined accordingly. Thus we obtain

²²In most cases, of course, programs are equipped with some form error handling that will prevent or catch infinite loops or memory faults. But even in that case, a suited interpretation and translation into a penalty or another constraint-handling response is not straightforward.




Figure 4.8: Stochastic sampling selection schemes. (a) Stochastic sampling with replacement, performed as N independent roulette wheel spins, where N is the number of individuals to be selected. (b) Stochastic universal sampling; Roulette wheel is spun once, simultaneously selecting N individuals at equidistant positions, where the probability value q can be interpreted as the spin angle.

a biased wheel. The simplest form of bias is obtained by sizing the segments directly proportionally to the individual's fitness. For the selection probability of an individual A , we hence have that,

$$p_A = \frac{\mu(A)}{\sum_{B \in \mathbf{P}} \mu(B)}, \quad (4.1)$$

where μ is the pay-off of an individual and \mathbf{P} denotes the population.

The following main drawbacks are associated with roulette wheel selection (De Jong, 1975):

(1) Individuals with a fitness substantially above average may quickly become dominant. Depending on the scaling of the fitness function, the selection probability of such individuals, often termed *super individuals*, may become a number only insignificantly smaller than one, consequently leaving only a small fraction of chance to the remaining solutions. This, in turn, seriously impedes exploration and leads to a premature convergence.

(2) In later stages of a GA run, the variance of fitness values within the population can be expected to drop. Again depending on the scaling of the fitness function, the roulette wheel may exhibit almost equally sized segments. As a consequence, the selection pressure is drastically reduced, and so the detection of hyperplanes that are to be mixed is impaired.

To rectify these issues, De Jong (1975) proposed an additional scaling of the fitness function, which takes the fitness distribution within the population into account. The following three scaling mechanisms are typically used (Goldberg, 1989, pp. 123, 124):

(1) *Linear scaling*: This type of scaling is aimed at a selection proportion that, on average, leads to one copy of the averagely fit individual in the population and a defined multiple of copies of the best solution.

(2) *Sigma scaling*: In order to decouple the fitness from the variance in the population and, in so doing, by equalizing the selection pressure throughout the entire GA run, Forrest (1985) developed a scaling that normalized the raw fitness by the standard deviation. A variant, called *sigma truncation* was introduced by Goldberg (1989).

(3) *Power scaling*: This straightforward scheme was originally demonstrated by Gillies (1985). He scales the raw fitness function with an exponent in order to adjust the spread of population members with regards to their fitness. This approach is best combined with a generation-varying determination

of the exponent: for example, using small exponents in the early GA stage (emphasizing diversity) and larger exponents later to increase the selection pressure.

4.5.2 Rank-based Selection

The development of a robust scaling routine with a negligible bias and a minimum deviation between observed and expected number of trials is often far from being straightforward. Especially in the regime of multi-objective optimization problems—problems with multiple, incommensurable objectives—or constrained problems, scaling may lead to a skewed fitness distribution and hence to an abysmal convergence behavior. Baker (1985) attempted to mitigate this problem by using a selection scheme in which the expected number of trials is determined by the rank of an individual in the population rather than its absolute fitness.




Figure 4.9: Linear-rank-based selection. Selection probability p_A as a function of the rank of individual A . A selection pressure s of one leads to an equal probability of all individuals; greater pressures lead to an accentuation of higher-ranked candidates.

Under *linear ranking*, the selection probability for an individual A is given as:

$$\frac{1}{N} \left(2 - s + 2(s - 1) \frac{R(A) - 1}{N - 1} \right), \quad (4.2)$$

where N is the population size, $R(A)$ is the rank of solution A in the population; where $R(A) = 1$ is the solution with the lowest fitness and $R(A) = N$ is the candidate with the highest fitness in the population. The parameter s denotes the selection pressure. Its effect is illustrated in Figure 4.9: A value of one causes all individuals to be selected with equal probabilities, thus imposes no selection pressure, whereas a value of two leads to a linearly increasing pressure as a function of the rank of the solution.

Other, *non-linear ranking* schemes, for example, to accentuate the expected number of trials of highly ranked individuals, have been devised by other authors (Michalewicz and Attia, 1994).

4.5.3 Stochastic Universal Sampling

Stochastically, GA populations consist of a small number of samples, such that selection inevitably induces sampling errors. Standard roulette wheel selection, for examples, leads to high variance (a

large spread of actually drawn instances compared with the expected values). Any bias associated with selection errors can have dramatic consequences for the course of the GA such as leading to a premature convergence or the loss of hyperplanes with an above-average fitness. The *stochastic universal sampling* technique proposed by Baker (1987) minimizes the variance of selection and, like standard roulette wheel selection, does not show any bias. It requires only one (normally distributed) random number for all selection events, hence giving a selection scheme without replacement. Given N , the size of the mating pool, or selection events, the random number is generated from the range $[0, 1/N]$, denoting the first event. The certain event is segmented into N aliquots, each of which determines one selection event. The procedure can be illustrated by considering a wheel of fortune that has N equally spaced instead of a single pointer (see Figure 4.8(b)). In addition to a bounded spread, this method is also computationally more efficient than standard roulette wheel selection. This property, however, is of minor importance in many real world optimization problems in which the fitness evaluation strongly dominates the overall computation time.

4.5.4 Deterministic Sampling and Remainder Stochastic Sampling

Deterministic and *remainder stochastic sampling* are fitness-proportional schemes in which the expected occurrences are computed as in roulette wheel selection (Brindle, 1981; Goldberg, 1989). The integer fraction of the expectations, however, directly determine the number of trials in the mating pool. In *deterministic sampling* the remaining slots of the mating pool are sequentially filled from a list of individuals that is sorted by the fractional part of the expectations. *Remainder stochastic sampling with replacement*, in contrast, takes the fractional portion as the segment sizes for a subsequent roulette wheel selection through which the mating pool is complemented. And *remainder stochastic sampling without replacement* sequentially traverses the fraction-sorted list. Each fraction is then taken as the success probability for the selection of the corresponding individual.

4.5.5 Tournament Selection

In contrast to the former selection schemes, *tournament selection* (Brindle, 1981) does not depend on the fitness statistics in the population from which individuals are selected. Instead it utilizes an intermediate selection pool. A number t of individuals is randomly chosen from the population. The individuals in this pool undergo a tournament, whose winner is selected to participate in mating. The tournament is usually performed in a straightforward manner: The individual with the best fitness wins it. A common value for t , the size of the intermediate selection pool, is two. Figure 4.10 demonstrate the steps involved in this case, termed *binary tournament selection*.




Figure 4.10: Tournament selection. (a) A selection pool is created by randomly choosing t (two, in this example) individuals; (b) the fitter of both solutions is finally selected.

More formally, the flow of tournament selection is shown in Algorithm 4.2, where S denotes the intermediate selection pool. There, a more general variant is presented in which the winner of the tournament is determined probabilistically ($p_s \in]0.5, 1]$, denoting the selection probability of the winner). As will be shown later, the role of this parameter is to adjust the selection pressure.

Algorithm 4.2: Tournament Selection

```

1  $S \leftarrow$  randomly choose  $t$  individuals as selection pool
2 sort  $S$  (descending fitness values)
3 forall the individuals  $A$  in  $S$  do
4    $p_s \leftarrow$  generate (uniform) random number
5   if  $p_s \leq p$  then
6     return selected individual  $A$ 
7 return least fit individual

```

When analyzing tournament selection, the two phases can be distinguished (Poli, 2005):

- (1) the random sampling of individuals and
- (2) the tournament.

Similarly to standard roulette wheel selection, the sampling phase induces a selection error that may, in the extreme case, lead to a situation in which the best individual is not selected at all. The probability of an individual to be considered during sampling for one tournament is given by:

$$p_S = 1 - \left(\frac{N-1}{N} \right)^t, \quad (4.3)$$

where N is then population size. Assuming a standard GA with a generational gap of one, the number of selection events to fill the next population is N . Consequently, the probability that sampling does not consider one individual is (for a derivation see Poli, 2005):

$$p_n = \left(\frac{N-1}{N} \right)^{tN}. \quad (4.4)$$

Thus the fraction of individuals not considered for selection in one generation increases with a decreasing tournament size.

On the other hand, a large tournament size leads to an increased selection pressure. Even the smallest possible tournament window of two results in an expectation that is approximately twice as large as that of the median solution. The selection probability for one individual to be selected can be estimated as follows (Bäck, 1994):²³

$$p_A = \frac{R(A)^t - (R(A) - 1)^t}{N^t}, \quad (4.5)$$

where $R(A)$ is the rank of solution A , with N being the best and 1 the worst solution. To allow smaller selection pressures, a probabilistic acceptance as shown in Algorithm 4.2, where the winner is selected with a predefined probability, can be employed. From (4.5), Motoki (2002) derives the expected number of non-selected individuals:

$$\mathcal{L}(N, t) = \frac{1}{N} \sum_{A \in \mathbf{P}} \left(1 - \frac{R(A)^t - (R(A) - 1)^t}{N^t} \right)^N, \quad (4.6)$$

where \mathbf{P} denotes the population. This loss of solutions increases with the tournament size and also with the number of individuals in the population, albeit decliningly. Motoki (2002) conjectured that with a tournament size of two, about 40 percent of the population will be lost due to tournament selection. Blickle and Thiele (1997) proposed to relate the loss of diversity to selection intensity, which measures the change of the average fitness under selection.²⁴ More precisely, they estimated the degree of diversity

²³The author makes the strict assumption that all fitness values in the population are different, or at least that a strict ranking order can be applied. A slightly different approach, in which instead of the rank the cumulative distribution function (CDF) of the fitness distribution was used was proposed by (Blickle and Thiele, 1995).

²⁴The notion of selection intensity has been transferred from quantitative genetics to genetic algorithms by Mühlenbein and Schlierkamp-Voosen (1993).

one can expect when aiming at a certain selection intensity. They found a similar behavior as for ranking selection. When applying the refined formula of (4.6), however, a less advantageous relation between selection intensity and diversity loss can be expected.

Sokolov and Whitley (2005) propose a variant of tournament selection in which, in the case of binary tournaments, each individual is drawn exactly twice. The first individual of each competition is drawn by its rank, and its competitor is drawn from a randomly shuffled version of the population. Tournaments consisting of two instances of the same solution are prohibited. This approach reduces the spread associated with random sampling phase of standard tournament selection.

Goldberg (1990) developed *Boltzmann tournament selection*. The basic idea is borrowed from simulated annealing: Individuals are selected by subjecting neighboring solutions to a tournament with expectations following the Boltzmann distributions. According to Goldberg (1990), this selection scheme serves two purposes: (1) A transfer of the asymptotic convergence behavior of simulated annealing to evolutionary algorithms, and, (2) by virtue of its local tournaments, it could be applied to niching techniques (see Section 4.9).

4.5.6 Truncation Selection

Mühlenbein and Schlierkamp-Voosen (1993) introduced a selection scheme that is often exercised by breeders and termed *truncation selection*. Only a fraction of the population, which contains the best solutions, is eligible for selection. Solutions are drawn (uniformly) randomly from this selection pool. Blickle and Thiele (1997) associated a high loss of diversity with this scheme, which can be minimized only at the expense of a reduced selection intensity.

4.5.7 Selection of an Adequate Selection Scheme

Despite a number of comprehensive surveys and mathematical models describing the effects of different selection schemes (for example, by Goldberg and Deb, 1991; Blickle and Thiele, 1997; Motoki, 2002), no conclusion can be drawn as to a one best selection scheme. On the one hand, this is because selection interacts with the other genetic operations such as recombination and mutation and can ultimately only be evaluated in this context. On the other hand, there is no optimization scheme superior to any other (cf. Wolpert and Macready, 1997), and this general premise also holds for selection.

Fitness proportionate schemes lead to a fitness-function-dependent growth ratio that is large in an early GA stage and progressively decreases with the fitness variance of the population. Without any counteracting measures, the high growth rate also renders this scheme prone to a premature loss of diversity. The second origin of diversity loss, the principally unlimited variance of actually selected solutions, can be drastically reduced using *stochastic universal sampling*. Scaling techniques can be employed to adapt selection pressure. Their implementation, however, is highly problem-dependent and may result in a bias in itself. Especially when the fitness function consists of multiple objectives or under constrained optimization problems, the scaling needs to be thoroughly adjusted. The selection pressure of *truncation selection* strongly depends on the threshold parameter, the ratio of selected individuals. It can be associated with a high diversity loss, which is strongly biased since genetic material in individuals with a fitness below average is completely lost. *Ranking selection* and *tournament selection* can be adjusted to exhibit similar growth ratios. Assuming equal selection intensities, *tournament selection* and *exponential ranking selection* exhibit a higher loss of diversity than *exponential ranking*. Since in tournament selection the loss of diversity depends both on the population size and the tournament size, a minimum tournament size of two can be considered expedient.

Simply speaking, selection can be associated with two contradicting effects: the growth ratio of solutions that exhibit a high fitness and a loss of diversity. A useful metric to express this trade-off is to evaluate the loss of diversity as a function of the selection intensity (Blickle and Thiele, 1997). A direct approach to moderate diversity loss is to adjust the selection pressure or an application of genetic operators such as high mutation rates in order to restore genetic material that has been abolished due to selection (Goldberg and Deb, 1991).

However, it is not diversity in general that a genetic algorithm depends upon but “useful” diversity—that is, genetic material that can be expected to contribute to the solution—that needs to be preserved. For example, alleles showing a negative impact on the fitness under all circumstances should be rather removed from the gene pool. Of course, in most situations the determination of the usefulness of certain diversity is not predictable. And a loss of alleles reduces the accessible search space, potentially preventing the GA from exploring beneficial regions. A number of methods specifically aimed at maintaining productive diversity have been devised. These methods are commonly termed *niche techniques* and are discussed in Section 4.9.

4.6 Recombination

In standard genetic algorithms, the information exchange between solutions plays a central role. Simply put, it creates new samples exhibiting beneficial contributions from different search regions. This is a stochastic process that involves a large number of iterations, samples and selection, the operation that actually draws the samples that are to be combined. From this perspective, that is, regarding recombination rather as a process than as a single operation, it can be considered as the determination and combination of (statistically) optimal hyperplanes. As was shown for the binary case in more detail in Section 4.3.3, the hyperplanes are given by the defining alleles of the equivalence class to which the corresponding individuals belong. This is more evident in the case of discrete alphabets but also applies to real-coded GAs (Radcliffe, 1991). Moreover, since each individual represents numerous such hyperplanes, recombination can be thought of as massively, implicitly parallel process.²⁵

Traditionally, genetic algorithms employ fixed-length chromosomes, where the alleles are from a discrete, finite alphabet. We will restrict our focus to binary-coded genetic algorithms, that is, an allele is either zero or one. Recombination in this context is then commonly performed as string operations and is referred to as *crossover*, following the natural role model of genetics. Two crossover operators—illustrated in Figure 4.11—have been commonly and successfully applied to variety of applications:²⁶ (1) *n-point crossover* and (2) *uniform crossover*. Both take a very different perspective on mixing of hyperplanes, which will be shortly reviewed in the following subsections. In addition, a number of specialized crossover operators have been devised, as already discussed in Section 4.3.

4.6.1 *n*-point Crossover

The idea of *n*-point crossover is loosely adopted from genetics (Holland, 1975). It intermixes segments of (typically) two parents to form two offspring solutions. The segments are defined by randomly obtained *crossover points* that can be thought of as fraction sites at which the parents are broken apart into $n + 1$ parents. One and two-point crossover are illustrated in Figure 4.11(a) and (b), respectively.

Let A and B be the parent individuals, A' and B' the offspring solutions, and l the fixed bit-string length. Further, let

$$(s_0, \dots, s_n) = ((0, \chi_1 - 1), (\chi_1, \chi_2 - 1), \dots, (\chi_n, l - 1))$$

²⁵The implicit parallelism of GAs has been subject to debate and intensive research, some of which will be discussed in Section 4.10.1.

²⁶A should be noted that the wide application of these operators does not imply a superior performance in comparison with other routines (Wolpert and Macready, 1997).




Figure 4.11: Common crossover schemes in binary-coded genetic algorithms. (a) *One-point crossover*: a crossover point is randomly chosen, dividing the parent individuals A and B into two segments. The first child, A' receives its left segment from A and the right one from B . B' is generated from the respective other segments. (b) In *two-point crossover*, the outer segments of parent A are copied to the offspring solution A' . The center part is transferred to B' . Vice versa, individual B contributes its outer segments to B' and the center segment to A' . (c) A gene-wise mixing is performed in *uniform crossover*. The alleles in A' are randomly either taken from A or B (and vice versa in B'). On average, each offspring solution hence receives half of its genetic material from either parent. However, the ratio can also be adjusted, for example, in order to control the degree of disruption.

be the segments defined by the (ordered) crossover points χ_k . Formally, we then have that,

$$a'_i := \begin{cases} a_i & (i \in s_j)(0 \leq j \leq n \wedge j \text{ even}) \\ b_i & (i \in s_j)(0 \leq j \leq n \wedge j \text{ odd}) \end{cases}$$

and

$$b'_i := \begin{cases} b_i & (i \in s_j)(0 \leq j \leq n \wedge j \text{ even}) \\ a_i & (i \in s_j)(0 \leq j \leq n \wedge j \text{ odd}), \end{cases}$$

where i is the locus ($i = 0, \dots, l - 1$) and a_i is the corresponding allele in individual A .

Original studies of genetic algorithms proposed one-point crossover (Holland, 1975). The reasoning behind this choice is that a minimum number of crossover points also minimizes the level of disruption of hyperplanes. However, one-point crossover also introduces a strong bias of the defining length and the position of the defining alleles of hyperplanes. For example, hyperplanes defined in the left or right most segments of solutions will be recombined only with an insignificant probability. Multi-point, especially two-point, crossover have been proposed to counteract this bias (De Jong, 1975; Goldberg, 1989; Eshelman et al., 1989; Spears and Jong, 1991).

4.6.2 Uniform Crossover

This type of crossover has been proposed by Syswerda (1989). It takes a slightly different perspective on mixing than the n -point crossover techniques in that it employs a gene-wise rather than a segment-wise recombination (illustrated in Figure 4.11c). Given two parent solutions A and B ; for each gene where A and B exhibit different alleles, the corresponding locus in the first offspring solution, A' , is randomly taken from either A or B . The second offspring, B' , receives the respective allele from the opposite parent. Originally, the probability was fixed to 0.5. That is, each offspring receives half of its alleles from either parent. However, any swap ratio can be applied.

Uniform crossover introduces a significantly higher level of disruption than one or two-point crossover. In other words, it leads to a much increased probability of the redemption of competitive hyperplanes. Following the schema theorem, this fact drastically impedes the exploitative efficiency of the GA and thus leads to a convergence degradation. However, at the same time it also increases the accessibility of new search space regions and so the exploratory power. In fact, Spears and De Jong (1991) have shown that especially when used with small population sizes, the disruption caused by uniform crossover may become productive. Moreover, they argue that the degree of disruption can be conveniently controlled by the swap ratio. In so doing the operator can be adapted to different situations such as the population size. Finally, the author argue that uniform crossover, in contrast to n -point crossover, does not introduce a bias with respect to the defining length (distance between the defining [fixed] positions) of hyperplanes.

A variant of uniform crossover, which is also aimed at alleviating the problem of positional and defining length bias, was devised by Caruana et al. (1989) and called shuffle crossover.

4.7 Mutation

In contrast to evolution strategy, which in their traditional form mainly relies on mutation, it was long regarded as a secondary, almost optional operator in genetic algorithms. Its main purpose was described as re-introducing genetic material that has been lost in the course of evolution (see, for example, Goldberg, 1989). In fact, an increase of diversity of candidate solutions in the GA population—in order to prevent a premature convergence toward a sub-optimum or a local optimum in highly multimodal search spaces—is an important function. Without mutation, the upper limit of possible alleles is determined by the randomly generated initial population. In other words, the GA's search space is restricted to those hyperplanes that are already present in the first generation. In addition, by preferring solutions that (currently) exhibit a pay-off value above average, selection eventually leads

to an allele depletion. Analogous to genetics, this phenomenon is termed *genetic drift*. By randomly introducing new alleles, diversity can be re-increased, allowing additional parts of the search space to be explored. In other situations, mutation may also be suited to increase the convergence rate, for example, by yielding well-performing individuals by fluke or by acting as a short-range variation scheme that accelerates the ascent (or descent) of solutions. As mentioned earlier (Section 4.3), in the scope of other than binary representation, a number of mutation operators have been devised, many of which play a central role in the evolutionary algorithms associated with them.²⁷ Another perspective when comparing mutation and recombination has been put forth by Spears (1993). He shows that the exploratory power of mutation is in principle unlimited since, in contrast to crossover, it does not “respect” any given hyperplane distribution in the population. In terms of construction, that is, in combining low-dimensional hyperplanes to form higher order ones, his findings indicate an advantage of crossover over mutation. Ultimately, the effect of genetic operations is strongly problem-dependent (Wolpert and Macready, 1997). Thus GA variants that are mainly driven by mutation have been proposed and are not categorically less efficient than their conventional counterparts.

The most common mutation operator for bit-string GAs is *point mutation* in the form of bit-flip operations (Holland, 1975). The chromosome is traversed, and for each gene a random number is drawn whose comparison with a predefined mutation probability determines whether to invert the allele. The random number is commonly uniformly distributed. But other distributions are applicable as well. Since normally a low mutation probability is chosen, only a few alleles of an individual are modified through mutation. Point mutation is illustrated in Figure 4.12.

```

      #
1001101110
1000101110

```

Figure 4.12: Fictitious genetic algorithm point mutation example. A low mutation probability is applied such that only a single allele is flipped.

4.8 Control Parameters

The behavior of the genetic operators such as crossover and mutation is controlled by parameters. The crossover rate (often written p_c), for example, determines the probability that the selected individuals mate. In general, in order for the evolution dynamics to take effect, this ratio is a value between 0.5 and 1. In other words, half to all of the selected individuals actually perform crossover. The remaining selected pairs are copied verbatim into the next population, mainly to increase the survival probability of candidates with an above-average fitness. The role of the mutation rate parameter (p_m) strongly depends on the mutation operator that is employed. For point mutation, it determines the probability of an allele to be inverted or, equivalently, the average inversion fraction of the chromosome. Normally, mutation rates are set to low values, in the order of 0.01 or below.²⁸ High mutation rates are generally considered obstructive for the hyperplane assembly associated with recombination—since they typically impose a random-walk-like behavior on the GA. A mutation rate close to zero, however, may impair exploration and so prematurely leave the algorithm in a steady state.

Like the mutation and crossover probabilities, the population size can be regarded as yet another control parameter. The overall runtime of an evolutionary optimization is typically determined by the

²⁷A discussion on this was, for example, given by Tate and Smith (1993).

²⁸Another common recommendation is $1/l$, where l is the length of the chromosome (for example, De Jong, 1975; Mühlenbein, 1992).

number of function evaluations, which in turn have an upper bound given by the number of solutions (per iteration) times the number of iterations. The population size is hence a critical factor for the efficiency of an evolutionary algorithm. On the other hand, it also determines its efficacy, by facilitating or impairing the underlying processes. A number of different theorems and recommendations concerning ideal population sizes have been presented: Goldberg (1989), for example, estimates the population size in terms of an optimal hyperplane processing. The minimization of sampling errors during the course of hyperplane mixing reveals a different picture and so a different estimate (Goldberg et al., 1992b); it was later supplemented by a building-block (hyperplane) supply and selection model (Harik et al., 1999). Many other studies are of empirical nature (e.g., De Jong, 1975; Schaffer et al., 1989) or use rigorous search-based approaches (Grefenstette, 1986) for the determination of a suited sizing. In Mühlenbein and Schlierkamp-Voosen (1993), the authors attempted to conjecture a semi-empirical rule, which suggest that the minimal population is a function of the problem size, the selection intensity (see Section 4.5) and the allele distribution in the initial population.²⁹

The selection of adequate control parameters may have a significant impact on the convergence behavior of evolutionary algorithms. Their determination is therefore of paramount importance. Many authors have tried to empirically identify ideal parameters (just a few examples are De Jong, 1975; Schaffer et al., 1989; Bäck, 1993). Since the GA control parameters are just as problem dependent as the operators themselves, it is little surprising that except for a number of qualitative statements there is no conclusive agreement between the results of the different studies. A more rigorous approach was presented by Grefenstette (1986). He applied a meta-search for optimal control parameters. The underlying optimizer was again a genetic algorithm (thus applying a meta-evolutionary approach). In his work, ideal parameter configurations drastically differ depending on whether the evaluation is performed on offline or online parameters.

In order to adjust GA parameters, numerical experiments are often required. This can, for example, be done by running the GA with different settings for a few iterations and then comparing the convergence behavior with a view to infer ideal parameters. Noting that GA control parameters may mutually influence each other (Eiben et al., 1999), the sheer number of possible combinations may already render this procedure infeasible, especially in the scope of numerically complex function evaluations. As an alternative, dynamic parameter control mechanisms have been suggested, which can be distinguished in various ways. Eiben et al. (1999), for example, present a meta-classification for several aspects of parameter control.³⁰ Here, we loosely follow their classification based on the underlying parameter control mechanisms:

- (1) *Deterministic/dynamic*: the control parameters take no feedback; parameters are either constants or their adjustments follow predefined rules, for example, as functions of time.
- (2) *Adaptive*: the control parameters are adapted to the state of the EA.
- (3) *Self-adaptive*: the control parameters themselves are subject to optimization, for example, by supplementing individuals with control parameters and exposing them to the same evolutionary process, or in a meta-evolutionary process.

As shown in Section 4.1.1, evolutionary programming and evolution strategy can be considered the earliest approaches to comprise a framework of different adaptive and self-adaptive parameter control schemes. It was not until the 1990s that similar strategies also gained popularity in genetic algorithms.³¹ In contrast to ES and EP, parameter and operator configurations have typically been considered as constants. Much of the research concerned with the impact of control parameters was consequently directed toward finding universally applicable configurations. And indeed, the success of GAs following the recommendations so obtained seems to perfectly justify this procedure. Originally, adaptive techniques have been applied to GA variants and problems with a strong sensitivity against parameter settings. There, parameter adaptation techniques are aimed at improving the convergence behavior in terms of (1) the best local optima that can be obtained, (2) robustness (reproducibility of results) and

²⁹The authors' benchmark was a OneMax problem (the maximization of ones in a bit string), which was chosen because the emphasis of that study was on additive objectives.

³⁰The same reference also contains a comprehensive survey on parameter control techniques; with later extensions in Eiben et al. (2007).

³¹Despite some early considerations, for example, by (Rosenberg, 1967).

(3) convergence speed. In addition and as a consequence, a successful adaptation technique can also be expected to increase usability, since it reduces the number of configuration settings that have to be determined and provided prior to the actual optimization run—at best leading to a genetic algorithm without parameters (see details, the study presented by Bäck et al. (2000a) and the brief survey in Appendix E). As discussed before, the GA convergence behavior can be coarsely characterized as a balance between exploration and exploitation. Adaptive methods can thus be expected successful if they have a beneficial impact on this tradeoff. In that sense, parameter adaption is one technique *inter alia* and competes, for instance, against niching methods (Srinivas and Patnaik, 1994).³²

Adaptive and self-adaptive approaches have proved successful in a variety of different applications, where in some cases they clearly outperform static operators and control parameters. The success can be attributed to two effects (Schwefel, 1987; Bäck, 2000; Spears, 1995):

- (1) *Diversity* of operators or control parameter settings. The mere existence of different configurations and their rotational application may lead to a convergence improvement, for example, by increasing the jump probability to different search regions or by preventing a pre-mature convergence to local optima.
- (2) *Selection of an optimum configuration*. The ideal of adaptive parameter control is to find an operator and parameter configuration that, in the given context, leads to an optimal optimization behavior of the genetic algorithm. Especially for self-adaptive methods, it is virtually impossible to ascertain whether or not self-adaptation is in action. Of course, empirical studies can be conducted comparing performances of static, adaptive and self-adaptive methods, however, not allowing to shed much light on the “adaptation power” of the method in question.

In order to draw general conclusions about the applicability and the performance of automated parameter control methods, further theoretical investigations are required. As a result from the NFL theorem (Wolpert and Macready, 1997), adaptive GAs are not generally superior to fixed parameter versions. Recent studies, however, have shown that champion-producing, co-evolutionary approaches can be associated with free lunches (Wolpert and Macready, 2005). That is, for certain co-evolutionary settings, some algorithms outperform others. Further investigations are hence required to clarify to what extent this result can be utilized in the scope of parameter adaptation.

4.9 Diversity Preservation

As mentioned before, the search behavior of evolutionary algorithms is often characterized as a balance between exploration of the search space and exploitation of the so gathered information in order to converge to local extrema.³³ An evolutionary algorithm putting too much emphasis on exploration is consequently prone to showing an impaired convergence performance. An algorithm neglecting exploration, on the other hand, tends to prematurely converge and so to yield a globally suboptimal solution. To counteract the latter behavior, increased mutation probabilities are often used with a view to prevent a genetic drift. Alternatively, dynamic or adaptive control parameters have been proposed as a means to increase the diversity of the individuals belonging to a population (Srinivas and Patnaik, 1994). The need for an intense exploration aggravates with the degree of modality of the search space. A number of techniques are therefore specifically aimed at maintaining solutions at multiple optima. Examples for such techniques are mating restrictions (for example, see Deb, 1989), in which only similar solutions are allowed to yield offspring, or multipopulation approaches (for example, see Cohoon et al., 1987; Mühlenbein, 1991), in which the evolution of each population can be biased toward different search space regions. Another class of approaches aimed at the sampling of the search space is termed niching. The common idea of these methods is to impact the distribution of individuals and to so modify the exploration behavior of the evolutionary process. Broadly speaking, a niche can be said to be a region

³²See Section 4.9 for a discussion on niching techniques.

³³This can be considered a rather superficial characterization and has been challenged by some authors (e.g., Eiben, 1998). Most importantly, the association of *exploitation* with selection and *exploration* with recombination and mutation proves untenable.

in the search space that exhibits a common local optimum. Given a search space D and a distance metric d ,

$$d : D \times D \rightarrow \mathbb{F},$$

where \mathbb{F} is an ordered field; a local maximizer can be defined as follows:

Definition 4.1 (local optimizer) $\hat{\mathbf{x}} \in D$ is called a local maximizer (without loss of generality (WLOG)): \iff

$$(\exists \delta \in \mathbb{F})(\forall \mathbf{x} \in D)(d(\hat{\mathbf{x}}, \mathbf{x}) < \delta \implies \mu(\hat{\mathbf{x}}) > \mu(\mathbf{x})).$$

Clearly, this definition allows to define an environment associated with the local optimizer, for example, by defining a connected, open neighborhood $N_{\hat{\mathbf{x}}}$ with

$$\mu(\mathbf{x}) \leq \mu(\hat{\mathbf{x}}) \quad (\forall \mathbf{x} \in N_{\hat{\mathbf{x}}}).$$

However, this definition does not allow to partition the search space in a straightforward manner since it includes overlapping niches. Another option would be to adopt the notion of *basins of attractions*, commonly used in local search approaches. It can be roughly described as a region or set of solutions that under a given operation, such as a local search operator, converge to the same attractor, for example, local optimum.

A pragmatic definition of niches, based on the position of solutions relative to local optima and hence ensuring that one solution belongs to exactly one niche can be put forward as follows: Let \mathcal{S} be a set of n locally optimal points, with

$$\mathcal{S} = \{\mathbf{s}_0, \mathbf{s}_1, \dots, \mathbf{s}_n\} \subset D.$$

Each optimal point spawns a niche containing all points to which it has a shorter distance than other optimal points:

$$\mathcal{N}_{\mathbf{s}_i} := \{\mathbf{x} \in D \mid \mathbf{s}_i = \arg \min_{\mathbf{s} \in \mathcal{S}} d(\mathbf{x}, \mathbf{s})\}.$$

The former definition does not enforce that the locally optimal points are optima with respect to the niche. Figure 4.13 shows such a situation in which an individual is close to a locally optimal point, which it outperforms. If not specifically accounted for, such niches are prone to extinction and can thus be considered insignificant. In his study on niche-preservation under *restricted tournament selection*—which we will introduce below—Harik (1995) omits such insignificant niches altogether by using the following optimal set definition:

$$\mathcal{S} := \{\mathbf{s} \in D \mid (\forall \mathbf{x} \in D)(\mathbf{s}' = \arg \min_{\mathbf{s}} d(\mathbf{x}, \mathbf{s}) \implies \mu(\mathbf{x}) \leq \mu(\mathbf{s}'))\}.$$

An explicit niching strategy, termed *sequential niching*, was proposed by Beasley et al. (1993). Their algorithm is specifically aimed at multimodal search spaces to which they apply multiple GA runs with adapting fitness functions. After one GA run yielded a local optimum, the respective region in the search space is blocked through a modification of the fitness function. This procedure is repeated until the intended number of optima are found. Other approaches include ecological GAs (Davidor, 1991), co-evolutionary sharing (Goldberg and Wang, 1998) and multi-national GAs (Ursem, 2000). The two major niching strategies can be broadly categorized as: (1) *sharing*, in which the exploration of the search space is altered through adaptively modify the selection pressure of individuals, and (2) *crowding*, in which replacement is altered to account for local characteristics of the search space.

4.9.1 Sharing

In sharing techniques, the objective value of individuals is modified in dependence of local, for example, distribution information of individuals. The main idea is to increase the selection pressure of individuals that exhibit a high similarity with many other individuals (Goldberg, Richardson 1987). Individuals




Figure 4.13: Niches can be defined by a neighborhood relation of solution to locally optimal points. However, this definition does not enforce the spawning local optima to be optimal with respect to the niche. Point \mathbf{x} , for example, clearly lies closer to the local optimum \mathbf{s}' than to \mathbf{s}_0 and would thus fall into the niche of the former. At the same time, it outperforms the spawning local optimum. Without specific measures such a niche can be expected to be prone to rapid extinction.

from sparsely sampled, yet highly paying off regions are consequently preferred over those situated in crowded areas. This is typically achieved by scaling the original fitness with a niche count measure,

$$\mu'(A) := \mu(A)/n(A). \quad (4.7)$$

The niche count (n) is then defined as the sum of a similarity measure between the individual and any other individual of the population:

$$n(A) := \sum_{B \in P} \text{sh}(d(A, B)),$$

where d is the distance between two candidates. The sharing metric (sh) is defined such that solutions contribute to the niche count depending on their distance from the target solution. In order to restrict the influence of the niche count to the vicinity of the evaluated solution, a threshold (σ), below which solutions are said to be dissimilar, is introduced. Formally, we have that,

$$\text{sh}(d(A, B)) := \begin{cases} 1 - (d(A, B)/\sigma)^\alpha & d(A, B) < \sigma \\ 0 & \text{otherwise,} \end{cases}$$

where the parameter α can be used to influence the decay of the niche measure with a decreasing distance.

The distance between two individuals can be defined in various ways. One approach is to employ a problem-dependent measure, based on the phenotype of individuals. In numerical optimization, for example, one can use the (Euclidean) distance between the parameter vectors. In structural or geometric optimization problems one could define a metric based on the shape or the position of individuals. Another common technique is to measure the distance of individuals in the genotypic domain. In binary-coded GAs, the Hamming distance (bit difference) constitutes an adequate measure. For real-coded representations any norm such as the Euclidean distance may be suited.

One of the critical aspects of sharing is the selection operator used in the underlying genetic algorithm. Simple roulette wheel selection, for example, may, because of its short take-over times, strongly degrade the stability of niches. Therefore Oei et al. (1991) propose the use of tournament selection. However, they showed that a straightforward combination of tournament selection and niching leads to a chaotic behavior, resulting in a (potentially vast) loss of niches. They identified the problem to be a timing one: Broadly speaking, sharing can be said to be an adjustment of the number of individuals per niche, proportionate to their fitness. Selection can be considered the mechanism that exercises this adjustment. Since no feedback about the distribution in the target population being generated is taken into account, over-adjustments and hence the complete loss of niches are likely to happen. Consequently, Oei et al. (1991) suggest to continuously update the shared fitness values by an evaluation based on the target rather than the current population.

Another loss of niches can be attributed to an ill-selected value for the sharing parameter σ . Strictly speaking, the method requires that the distance between the local optima of a search space are known a priori. Moreover, since traditionally the same σ applies to all niches, only near-equidistant niches can be accounted for reliably (Deb and Goldberg, 1989). Additionally, in some respects sharing is similar to a bi-objective optimization problem, where one criterion is the actual objective whereas the other one is the sharing component. This renders sharing highly sensitive against ill-scaled fitness functions. To alleviate this issue, a problem-specific fitness scaling is in order (ibid.).

A variant of sharing was proposed by Petrowski (1996). In his *clearing* approach, a niche has a predefined capacity of individuals it can hold, resembling a competition for limited resources in that niche. Smith et al. (1993) provided a niching strategy in the context of artificial immune systems: A data point (antigen) is randomly selected. In addition, a predefined number (σ) of candidate solutions are selected at random. The solution candidate with the smallest distance to the data point receives a fitness appreciation. The rest of the σ individuals are unaltered. In contrast to fitness sharing, this method does not require any knowledge of the location of sub-optima in the search space. It is, on the other hand, only suited if, like in the presented antigen/antibody matching case, designated data points are known beforehand.

Clustering

In order to alleviate the time complexity required to evaluate the sharing component of the fitness value, Yin and Gernay (1993) propose a clustering approach, in which the search space is subdivided into niches prior to fitness evaluations. This technique reduces the number of individuals that have to be compared to the predefined partitions. Another clustering method was introduced Hocaoglu and Sanderson (1997). In their approach, genetic algorithms are independently run on multiple subpopulations, each of which is associated with one niche. After a predefined number of generations, the subpopulations are re-merged, and then distributed to new subpopulations, which are determined through a clustering method. Other clustering procedures have, for example, been proposed Miller and Shaw (1996); Gan and Warwick (1999); Streichert et al. (2004); Ando et al. (2005).

4.9.2 Crowding

Basic sharing techniques can be considered implicit niching techniques as they leave the control of the population structure to selection. An alternative approach is to modify the individual generation and replacement strategies in order for the population to exhibit a productive diversity. De Jong (1975) proposed a straightforward crowding method in conjunction with a steady-state GA. In his algorithm, newly created offspring solutions replace those individuals that they resemble closest. Resemblance can be defined in the same way as in fitness sharing, either genotypic or phenotypic. Although this approach is able to maintain a certain level of diversity, it is considered unsuited for highly multimodal optimization problems (Deb and Goldberg, 1989; Mahfoud, 1992). The latter author attributes these issues to a premature genetic drift due to stochastic errors in the replacement routine. He proposed a slightly different formulation, termed *deterministic crowding*, which he showed to outperform the original work of De Jong (1975) and related, derived approaches (see Mahfoud, 1995), especially in terms of replacement errors. It can be summarized as follows: The current population is traversed by randomly (in the original work, without replacement) drawing pairs of solutions. Let P_1, P_2 be such a pair. Conventional crossover and mutation operations are performed on each of the selected pairs, yielding two offspring solutions (C_1, C_2). Afterward, a local tournament between parents and children is performed for a spot in the next population. The pairing for the tournaments (T_1, T_2) is determined by the minimum average distance of parents and children,

$$T_1, T_2 := \begin{cases} (P_1, C_1), (P_2, C_2) & d(P_1, C_1) + d(P_2, C_2) < d(P_1, C_2) + d(P_2, C_1) \\ (P_1, C_2), (P_2, C_1) & \text{otherwise.} \end{cases}$$

This micro-selection ensures that each of the two offspring solution competes against the parent that it is most similar to. Thus the new population consists of a combination of individuals of the previous generation and akin, yet “fitter” individuals. In so doing, this procedure is aimed at maintaining a previously achieved level of diversity while, at the same time, converging toward optima in the similarity regions. Like before, distance can be defined either on the genotype or the phenotype, with the former being the typical options.

A variant, incorporating a probabilistic replacement scheme, was introduced by Mengshoel and Goldberg (1999) and termed probabilistic crowding. The probability of an individual \mathbf{x} winning a tournament against \mathbf{y} is taken as:

$$p_{\mathbf{x}} := \frac{\mu(\mathbf{x})}{\mu(\mathbf{x}) + \mu(\mathbf{y})}.$$

This is in contrast to deterministic crowding, where the individual with the higher fitness directly won the tournament. Mengshoel showed that the probabilistic component of his approach leads to the following proportion of a niche j :

$$N_j = \frac{\mu_j}{\sum_{i=1}^q \rho_{ji} \mu_i}, \quad (4.8)$$

where μ_i denotes the fitness associated with niche i , for example, taken as its average, ρ_{ji} is the transmission probability of individuals from niche j to niche i , and N_j is hence the proportion of

individuals in the j th niche. Equation (4.8) can be regarded as a general form of the *nicheing rule*, previously put forward by Mahfoud (1995).

Another crowding strategy was presented by Suzuki and Kobayashi (2005). It addresses two problems that the authors consider key to crowding and sharing approaches: (1) different sensitivities and inter-correlations of the optimization parameters, also depending on the search space region, which may cause genotypic and phenotypic distance metrics to fail; and (2) a nonuniform distribution of local optima, with dense occurrences in some search space regions. The main difference of this method is consequently the distance metric, which is defined as the average fitness along the path between two individuals. In other words, the higher the path fitness between two individuals, the higher is the probability that the two individuals are close. The authors derive this assumption from their definition of attractors, which they define as the convex about a local optimum (see Figure 4.13). This strict definition is, of course, applicable only to continuous search functions. The path fitness is obtained by samples, which are in turn generated through the same crossover operator as employed in their mating step. With the exception of the distance metric and a multi-parent mating scheme that is due to the employed crossover operator, the remainder of this method is similar to other crowding techniques.

Recently, Mengshoel and Goldberg (2008) presented an analysis framework for a wide range of numerous crowding techniques in which they show that many of these techniques drastically change the behavior of evolutionary algorithms. In fact, they claim that EAs making use of crowding are better described as *local tournament* algorithms and exhibit more similarities with approaches such as simulated annealing. This classification is also in accordance with the findings of Mahfoud and Goldberg (1995) and their effort in formulating a parallel, population-based simulated annealing approach, which adopts genetic operations from GAs but a dynamic replacement strategy that resembles the one of simulated annealing.

Restricted tournament selection

One approach that directly embodies the idea of localized tournaments was proposed by Harik (1995).³⁴ His technique utilizes the fact that standard tournament selection can be extended to account for other criteria in addition to just the objective value of individuals. Consequently, he altered the tournament such that only similar individuals compete against each other. The effect can hence be characterized as a competition for similarity slots in the population. Like in other niching approaches, similarity can be defined in various ways; specifically phenotypic or genotypic. The most common approach in restricted tournament selection (RTS) is to use a genotypic distance metric such as the Hamming distance in the case of binary coded GAs.

The procedure consisting of crossover, mutation and selection is depicted in Algorithm 4.3 and can be summarized as follows: Two individuals are randomly selected from the population. Two offspring are created using conventional crossover and mutation operations. For each of the offspring, the appropriate slot is determined. This is done by searching the current population for a solution that most closely resembles the child solution. Commonly, the search is restricted to a predefined number of randomly selected individuals (window). After the appropriate competitor has been found, a tournament is performed. That is, depending on the results of the fitness comparison, the original solution is maintained or replaced with the newly generated offspring.

Harik showed that restricted tournament selection exhibits a take-over rate that is in the same order as tournament selection, albeit, with the take-over restricted to individual niches.

Harik proposes a selection for the size of the window that is proportionate to the number of niches that are to be maintained, which in turn correspond to the level of modality. Selecting too small a size will decrease the take-over time of highly “fit” individuals and thus lead to a pre-mature loss of representatives in locally optimal regions. Too large a value, on the other hand, may lead to too fine a partition of the search space, which in turn may lead to an impaired selection pressure and thus to a slow convergence. With an adequate window size, restricted tournament selection has been shown suited to maintain niches for a large number of generations (Harik, 1995).

³⁴Later, an almost identical approach was proposed by Grüninger and Wallace (1996) and termed struggle GA.

Algorithm 4.3: Restricted tournament selection

```

Data:  $P^{(k)}$  // current population
Data:  $P^{(k+1)}$  // next population
Data:  $w$  // tournament selection window
1  $A \leftarrow$  randomly select from  $P^{(k)}$ 
2  $B \leftarrow$  randomly select from  $P^{(k)}$ 
3  $A', B' \leftarrow \text{crossover}(A, B)$ 
4 foreach  $C \in \{A', B'\}$  do
5    $C \leftarrow \text{mutate}(C)$ 
   // compile selection window
6    $W \leftarrow \emptyset$ 
7   for  $i = 1, \dots, w$  do
8      $D \leftarrow$  randomly select from  $P^{(k)} \setminus W$ 
9      $W \leftarrow W \cup \{D\}$ 
   // find solution in window that is most similar to offspring
10   $C' \leftarrow \arg \min_{C' \in W} d(C, C')$ 
   // binary tournament between offspring and old, similar solution
11  if  $\mu(C') > \mu(C)$  then
12     $P^{(k+1)} \leftarrow P^{(k+1)} \cup \{C'\}$ 
13  else
14     $P^{(k+1)} \leftarrow P^{(k+1)} \cup \{C\}$ 

```

A very similar approach based on a steady-state GA was introduced by Cedeño et al. (1994). Like in restricted tournament selection, their multi-niche genetic crowding (MNC) routine incorporates crowding aspects both in the selection and the replacement schemes. The first mating candidate is selected randomly (or sequentially). For its partner, a number of individuals are randomly drawn from the current population, yielding a preselection window. From this window, the individual with the shortest (phenotypic, in their example) distance to the first mating partner is selected. Crossover and mutation are performed yielding two (or one, depending on crossover implementation) offspring solutions. Steady-state replacement is then performed by first forming intermediate sub-populations, termed *crowding factor groups*, containing a predefined number of individuals. From each of these islands, the one most similar (again, phenotypically) with the offspring solution is preselected. The offspring finally replaces that preselected individual with the lowest fitness. The authors claim that their approach, applied to a DNA sequencing problem, succeeded in (a) preserving diversity, (b) generation stable niches and (c) foster the inner-niche convergence. The main difference compared with RTS is—in addition to the steady-state nature of their procedure and the use of a phenotypic rather than a genotypic distance metric—the fact that also through selection similar individuals receive a higher mating probability. The major caveat of this approach is the selection of suited parameters: the preselection window, the number of *crowding factor groups* and the number of individuals in each of the former.

4.10 Genetic Algorithms Theory

Numerous attempts have been made to analyze the behavior of genetic algorithms in order to derive theorems capturing their working principles, to formally prove their convergence or to assess their performance on specific problems. In this section, the classical schema theorem, its extensions and state-of-the-art model frameworks are briefly discussed.

4.10.1 The Schema Theorem

The first attempt to develop a theoretical framework of the schema³⁵ processing in genetic algorithms was exercised by Holland (1975) and his coworkers. Their *schema theorem* can be considered both a coarse estimate of the effect of genetic operators on schemata and thus also a reasoning for the specific design of genetic algorithms.

Assuming one-point crossover, point mutation and fitness-proportionate selection, the growth of a schema is lower-bounded by the following expression:

$$m(H, t+1) \geq m(H, t) \frac{\mu(H, t)}{\hat{\mu}(t)} \left(1 - p_c \cdot \frac{\delta(H)}{l-1}\right) (1 - p_m)^{o(H)}, \quad (4.9)$$

where m is the expected number of samples of schema H at a given generation, $\hat{\mu}(t)$ is the average pay-off of that schema at generation t , and as before, p_c and p_m are the crossover and mutation probabilities, respectively. The length of the binary chromosomes is denoted by l . The distance between the first and the last defining positions of a schema is called the defining length and written as $\delta(H)$. The order of schema ($o(H)$) is given by the number of defining bits. For example, given a chromosome of length (l) eight and a schema ****1**0****. The defining length ($\delta(H)$) is four, whereas the order ($o(H)$) is two.

Before discussing conclusion that can be drawn from (4.9), let us briefly outline its derivation (chiefly following Goldberg, 1989, pp. 28–41):

Selection

Let \mathbf{P}_t be the current population and $\hat{\mu}(t)$ its average pay-off,

$$\hat{\mu}(t) = \sum_{A \in \mathbf{P}_t} \frac{\mu(A)}{|\mathbf{P}_t|}.$$

Furthermore, let $m(H, t)$ be the number of instances of the schema H at time t . Under fitness proportional selection, the number of instances of the schema H in the next generation is hence:

$$m(H, t+1) = m(H, t) \frac{\mu(H, t)}{\hat{\mu}(t)}. \quad (4.10)$$

A straightforward conclusion is that fitness proportional selection assigns an increasing number of trials to schemata with an above-average pay-off. Let H be such an above-average-fit schema. Further, let c be a constant by which the schema is above average, that is,

$$\mu(H, t) = (1 + c) \hat{\mu}(t).$$

This allows us to rewrite (4.10) as:

$$m(H, t+1) = (1 + c) m(H, t).$$

If we now assume a c to be a constant over a number of time steps, starting with some initial time 0, i.e., the schema H stays above-average by some factor for some time, we obtain the following exponential growth rate:

$$m(H, t) = m(H, 0) \cdot (1 + c)^t.$$

Holland (1975) and De Jong (1975) compared this behavior to a game theoretical example, the *k-arm bandit*, where the goal is to maximize winnings at a slot machines after a number of games. The arms are associated with different average pay-offs with a certain variance, and one arm is known to have a higher expected pay-off than the others. The gambler, however, does not know which one of the arms performs best. Like in GAs, this problem poses a dilemma between exploration (finding the best performing

³⁵see Section 4.3.3

arm) and exploitation (pulling the best-observed arm to pocket its pay-off). One successful strategy for gamblers to circumvent this dilemma is to allocate slightly more than exponentially increasing trials to the arm currently observed best. Holland argues that this is close to identical to the principle and the effect of selection in genetic algorithms.

Genetic algorithms, however, employ operators, specifically recombination and mutation. In the context of schema growth (or decay), these operators can be said to have a disruptive impact on schemata. In other words they decrease their survival probability.

Recombination

When considering the survival probability of a schema, recombination is to be examined as disruptive force, that is, as an operator that potentially destroys a given schema. Under one point crossover, a schema is disrupted if the crossover point is situated at a locus within its defining length $\delta(H)$,

$$p_d = \frac{\delta(H)}{l-1}.$$

Considering, but not quantifying, situations in which offspring solutions belonging to H can also be created from other sources than from parents that themselves are from H and accounting for the crossover probability p_c , we obtain the following lower bound for the survival probability of schema H :

$$p_s \geq 1 - p_c \frac{\delta(H)}{l-1}.$$

Mutation

Point mutation stochastically independently acts upon the bits of a chromosome. Given the mutation probability p_m , the probability for a certain gene not to be mutated is $(1 - p_m)$. Thus for a schema with order $o(H)$ the probability not to be disrupted is:

$$(1 - p_m)^{o(H)}.$$

One of the conclusions Holland (1975) drew from the schema theorem was that low-order schemata, with a short defining length play an important role in genetic algorithm. Because of their lower disruption probability under genetic operations and without regarding other specifics of the genetic operators, their survival probability and thus their growth rate can be expected to be high. As will be shown in Section 4.10.4, the former statement can be considered an overinterpretation, resulting from the fact that Holland's schema theorem (4.9) does not account for the reconstruction of schemata through crossover or mutation. In fact, it does not make any statement about the effect of mixing on the evolution of schemata (other than the said lower bound). Furthermore, the schema theorem only provides a static, single-generation perspective and is hence not suited to describe the dynamics associated with the evolutionary process. This has led to a number of serious critiques about Holland's schema theorem (for example, Vose, 1993; Vose and Wright, 2001; Antonisse, 1989; Grefenstette and Baker, 1989; Mühlenbein, 1991; Radcliffe, 1991, 1992), challenging its usefulness and some of the conclusions drawn from it. One of the major concerns is the interpretation of "implicit parallelism" offered by Holland. Although the statement that solutions are simultaneously representing a number of different schemata is undoubtedly correct, any conclusion that this provides an advantage of genetic algorithms over other search routines is misplaced (Wolpert and Macready, 1997; Vose and Wright, 2001; Wright et al., 2003).

4.10.2 Computational Complexity

Critical schemata³⁶ can be considered as building blocks that are combined to form above-average candidate solutions. In the *building block hypothesis*, this conjecture was translated into an explanation on how genetic algorithms work:

³⁶Short, low order, high-performance schemata according to Goldberg (1989). We shall see later that this definition has been challenged by others (for example, Stephens et al., 2000).

“Just as a child creates magnificent fortresses through the arrangement of simple blocks of wood, so does the genetic algorithm seek near optimal performance through the juxtaposition of short, low order, high-performance schemata, or building blocks.” Goldberg (1989, p. 41)

As a consequence of this hypothesis, Bethke (1981) and Goldberg (1987) identified a category of problems they considered hard for genetic algorithms and which they refer to as deceptive problems. In short, if two building blocks are separated on the chromosome by a large distance, their combination is subject to a high disruption probability under recombination. If additionally the schema constituted by the recombination of these building blocks performs worse than, say, the recombination of the (bit-wise) inverse building blocks, the problem is said to be deceptive. Because disruption due to recombination will allow only a small number of the combined schemata to survive, and because of the “local” optimality of the individual building blocks, the GA shows an impaired convergence toward the combined, globally optimal schema.³⁷ In the extreme, a so defined problem of deception can be considered one in which a globally optimal site is surrounded by a suboptimal environment. In fact such a situation poses a difficulty for many other optimization approaches. Strictly speaking, however, this specific “needle-in-a-haystack” situation is directly linked to the way decision variables are represented in genetic algorithms. Consequently, Goldberg proposes a representation-related approach to circumventing deception. He argues that if the distance between deceptive building blocks could be reduced, the disruption of the combined larger schema is drastically decreased. In turn, the survival probability of the more successful larger scheme not consisting of the deceptive building blocks is increased, and the deceptive building blocks eventually vanish. Some of the so motivated approaches are summarized in Section 4.10.3.

Similarly to the derivation of deceptive search functions, different fitness landscape analyses have been performed in order to classify search spaces as “GA-hard” or “GA-easy” (for example, Tanese, 1989; Hart and Belew, 1991; Forrest and Mitchell, 1993b,a). None of the results could confirm a strict correspondence between GA-hardness and deception. As a consequence, Liepins and Vose (1990), for example, conjectured that deception is to be considered one of many factors that render a problem GA-hard or easy. Later, however, also this conception was seriously challenged by showing that equivalent search functions, of which one was deceptive and one was not, posed the same hardness for a GA (Grefenstette, 1993; Rowe and East, 1995).

Like the authors before, Davidor (1990) associates “GA-hardness” with genotypic properties, arriving at an epistasis concept for artificial evolution, which as in genetics can be broadly characterized as non-linear intercorrelations of genes. In order to analyze and optimize coding-induced issues, Davidor proposes a framework to estimate the epistasis variance, which reflects the “degree of non-linearity” of the phenotypic/genotypic relation of a problem. For a discussion of the state-of-the-art of epistasis research in GAs see Iglesias et al. (2005, 2008), for example.

Advocates of genetic algorithms have sometimes claimed that “their” approach exhibits a principal advantage over related approaches. This bold statement is resulted from a great number of successful applications but also from theoretical misconceptions, mainly associated with the notion of implicit parallelism (Holland, 1975; Goldberg et al., 1992b). Numerous studies on test functions and also several real world applications, however, clearly showed that GAs are not the general purpose routines they are sometimes considered. To this end, Radcliffe and Surry (1995) raised questions about performance and limitations of GAs in a general manner. One of the main conclusions drawn is that the problem representation plays a more important role than it was generally believed before. Consequently, the authors criticize the then common practice of assessing the performance of search and optimization methods on the basis of static benchmark functions. Having dealt with similar questions in the scope of machine learning, Wolpert and Macready (1997) formally proved that evolutionary algorithms do not exhibit any critical advantage over any other optimizer. Broadly put, their NFL theorems state that on average over all search spaces, all optimization algorithms perform equally well. In other words or as a practical conclusion, there is no intrinsic argument to choose evolutionary algorithms over other optimization methods. Culberson (1998), for example, argues that “*due to the complicated interaction*

³⁷Goldberg (1987) demonstrated this phenomenon by introducing what he called a *minimal deceptive problem*.

within an EA” (ibid.), other search methods, which obfuscate their inner workings less, may often be better suited. This statement, however, is prone to “violating” NFL by going to the other extreme, in that it ascribes an innate disadvantage to EAs. Consequently, Culberson acknowledges that evolutionary algorithms may show some advantages when applied to specific problems.

4.10.3 Linkage-aware Genetic Algorithms

Long before the building block hypothesis was devised, Bagley (1967) proposed the use of a coding that separates between alleles and loci of genes. This approach allowed him to employ an inversion operator by which the loci of genes are changed. His approach can be said to be implicit as no specific effort was taken to identify non-linearities between building blocks. Later, a number of explicit approaches were developed. In the *messy genetic algorithm* (Goldberg et al., 1989), a deterministic enumeration is employed in which all building blocks up to a predefined order—ideally including the maximum deceptive building block order—are generated. In a *primordial phase*, selection alone is applied to increase the proportion of above-average fitness values and to reduce the population size to a computationally manageable size. The subsequent *juxtapositional phase* proceeds identically to traditional GAs by selecting and mixing individuals that pay off well. The messy coding of genes is taken into account by a cut-and-splice recombination operator, which allows for variable bit-string length. Overspecifications of genes due to this operator and the messy coding are resolved by expressing only the first occurrence of a gene; underspecifications are repaired by employing templates, which are created from locally optimal candidate solutions. In contrast to the deterministic enumeration in the original implementation, the *fast messy genetic algorithm* uses what is called a probabilistic complete enumeration of building blocks (Goldberg et al., 1993a). The idea is to initialize the population such that it contains the intended order of building blocks with high probability.

Kargupta (1996) proposed a slightly different approach in his *gene expression genetic algorithm*. There, each gene is supplemented with an indicator denoting its impact on the fitness. A bit-flip test is performed on all individuals: If the fitness of the tested candidate increases due to the bit modification, the original allele is said not to belong to the class of best performing structures. The indicator of the gene is set to zero. Otherwise, if the fitness decreases under the modification, the allele is regarded as a contribution to above-average building blocks. The fitness difference between original fitness and fitness due to the bit-flip is stored in the indicator variable. Like in the related approaches, the indicator assignments are performed in a separate primordial phase, followed by a juxtapositional stage that operates similarly to conventional GAs. In order to assign more trials to above-average building blocks, a tailored operation, called class selection, but actually better described as a mixing routine, is performed: Two candidates are randomly picked. The offspring solution consists of the genes that exhibit a higher indicator value. To further advance mixing, class selection is supplemented with standard binary tournament selection. Recombination is implemented as a modified two-parent uniform crossover, in which a locus is randomly selected; if its weight is larger than the corresponding one in the other mating partner, alleles are exchanged.

Due to the building block enumerations in their primordial stages, the above algorithms may amount to a drastically higher computation time. The authors argue that the time spent during the enumeration pays off, since it allows the actual GA, that is, the juxtapositional phase, to converge significantly faster, or even find solutions where traditional GAs do not. It then depends on the nature of the problem—if it is assumed to be deceptive or not—whether the application of messy GAs is sensible or not.

A different approach to account for linkage was taken with the linkage-learning genetic algorithm of Harik (1997). Instead of strings, chromosomes are there encoded as rings. Two-parent recombination is performed by defining one of the parents as the donor and the second one as the recipient. From the donor, a randomly defined chromosome segment is introduced into a random position of the chromosome of the recipient. The obtained offspring solution then comprises an overspecified chromosome. Overspecifications are resolved, through a gene expression phase, in a first-come-first-served manner by starting from a random gene and, in predefined or random direction, removing redundant genes. Harik was able to show that, in combination with the proposed recombination operator, the ring representation leads to a progressive increase of linkage of both well performing and deceptive building blocks.

He also found, however, that under traditional selection operators the linkage learning mechanism is too slow to catch up with a convergence toward local optima. Hence he suggested the use of restricted tournament selection (see Section 4.9.2) or, as an alternative, a probabilistic expression mechanism in which extinct alleles are re-introduced with some probability. The linkage learning algorithm has been demonstrated to perform well on a number of—albeit rather academic—examples.

4.10.4 Exact Schema Theorems

The original schema theorem (4.9) does not constitute an exact evolution equation but a lower bound for the growth of above-average schemata. This is because Holland (1975) accounted for the disruption probabilities under crossover but not for cases when schemata are reconstructed. Several attempts have been made to rectify this shortcoming and to devise an exact schema theorem.

Although not restating the schema theorem, a number of seminal facet-wise investigations have been presented, which have clarified the behavior of crossover on schemata, especially in light of schema reconstruction (Booker, 1987; Eshelman et al., 1989; Syswerda, 1989; De Jong and Spears, 1992). Bridges and Goldberg (1987) derived a formalism allowing to evaluate the exact number of copies of candidates under selection and (one-point) crossover. According to the authors, a translation to estimating the expected number of schema representatives is straightforward. Later, Whitley (1992) extended this work in that direction, albeit using a different model.

Goldberg et al. (1992a) emphasized the role of the population size. They devised a straightforward formulation to account for the impact of sampling biases on genetic operations. According to their estimates, one can find feasible population sizes that produce a negligible bias compared to the limit of bias-free infinite populations.

As discussed in Section 4.3, Radcliffe (1991)—and similarly Vose (1991)—developed a generalization of the schema theorem suggesting that the equivalence class processing of genetic algorithms is not restricted to binary or discrete alphabets of higher cardinality.

Altenberg (1995) proposed the application of Price's covariance and selection theorem as a performance measure for GAs. He showed that by incorporating suited measurement functions, the fitness distribution in one generation can be expressed as a function of the fitness distribution in the previous generation. The use of a different measurement, called schema indicator function, allows to directly derive the schema theorem. Thus the schema theorem can be considered a specialization of Prices's theorems. In addition, Altenberg introduced what he refers to the *missing schema theorem*, which describes the evolution of the fitness distribution under selection and recombination.

In their attempt to identify the effective degrees of freedom, Stephens and Waelbroeck (1998) propose to “coarse-grain” from the bit-level of to an average perspective, eventually arriving at an exact schema equation, which takes into account disruption and reconstruction of schemata. In addition, the authors introduce the concept of an effective fitness that accounts for fitness distributions under genetic operations. One of the remarkable conclusions the authors drew from their exact equation is that, in contrast to the building block hypothesis, there is no general preference of short, low-order schemata over longer ones (also see Stephens et al., 2000).

As demonstrated by Wright (1999) and Vose (1999, chap. 19), the dynamics system model by Vose et al. (explained later in Section 4.10.6) can be used to derive equivalent, exact formulations to compute the quantitative effect of recombination and mutation on schemata. The analysis is performed in the Walsh domain, where the scaling under selection, mutation and crossover is directly related to Walsh coefficients.

4.10.5 Dynamic Models

Albeit specifying the original schema theorem by including schema reconstruction through recombination and mutation, the former models are not suited to account for the dynamics of genetic algorithms. Wright (1999) argues that the main limitation on doing so is posed by selection. Different models have been derived to analyze the effect of selection on the fitness distribution, leading to different measures such as the growth rate and the takeover time of candidates (Goldberg and Deb, 1991), or

the selection intensity (Mühlenbein and Schlierkamp-Voosen, 1993). The latter metric, borrowed from the field of quantitative genetics, has been used by a number of other authors to study and compare different selection schemes (for example, Thierens and Goldberg, 1994; Bäck, 1995; Blickle and Thiele, 1997; Motoki, 2002). In order to assess the dynamic behavior of building block mixing, Goldberg et al. (1993b); Thierens (1996) propose a coarse-grained model in which they use a dimensional analysis ansatz to interrelate a simplistic expected mixing time estimate with their previously derived takeover time analysis.³⁸

In addition to these and other approaches, which could be classified as *facet-wise* studies on the GA dynamics, the following more rigorous, closed models have been proposed and devised.

Markov models

Evolutionary algorithms can be described as (time-discrete) dynamic stochastic processes. It thus stands to reason that a number of authors have proposed according models with which the dynamics of EAs can be analyzed. A common approach is to employ (finite) Markov chain models. There, each realizable population denotes a state, the genetic operations leading from one population to another are the transitions, which are associated with different transition probabilities, and a GA run can consequently be characterized as a certain trajectory through the states (populations). The Markov model is applicable since the transition from the current population to the next depends only on the current population (and a number of stochastic events) but not on any prior state. Simple (or canonical) GAs can moreover be modeled as time-homogeneous Markov chains. Goldberg and Segrest (1987), for example, have applied this technique to study genetic drift. Eiben et al. (1991) and Rudolph (1994, 1996) used it for convergence analyses. Specifically, they showed that GAs converge to global optima only under certain conditions. One of the significant findings was that EAs using elitism can be proved to converge to global optima. Davis (1991) devised a more complete Markov chain framework, in which he also considered time-inhomogeneous GA variants, for example, to account for time-varying mutation probabilities. Davis showed that even for time-inhomogeneous mutation the optimum individual cannot be guaranteed to stay within the population; in other words, the EA does not converge.

4.10.6 Dynamics Systems Models

A rigorous dynamic systems model has been proposed by Vose et. al (Vose and Liepins, 1991; Vose, 1999). It was originally based on the assumption of infinitely large population but later extended to account for finite populations (Nix and Vose, 1992). The main ideas and properties are shortly summarized as follows: Populations (states) are represented by vectors \mathbf{p} , whose components correspond to an enumeration of all possible points in the search space. The values of the components denote the occurrence of one such point in the population. That is, given a population size of r and k instances of a candidate A_i , we obtain a value of k/r for the corresponding state vector component p_i . The vectors \mathbf{p} form a simplex,

$$\Lambda = \{\mathbf{p} \mid \mathbf{1}^T \mathbf{p} = 1 \wedge p_i \geq 0\},$$

where, in the finite population case, only those elements of Λ whose components are rational numbers with a denominator of r are realizable populations. Another interpretation of elements in Λ is that they correspond to a sampling distribution. In other words, each p_i denotes the sampling probability of the corresponding candidate solution. The transition from one population to another can then be regarded as a function τ , with

$$\tau : \Lambda \rightarrow \Lambda.$$

In order to account for the probabilistic nature of GAs, the transition function is assumed to be a stochastic. This can be realized by regarding τ as an implicit rule whose explicit form consists of two steps: (1) the generation of a distribution assigning a probability to each candidate solution and (2) a

³⁸A comprehensive survey on the role of building block mixing and according models was offered by Goldberg (2002).

stochastic drawing of trials according to the former distribution. The according heuristic function that generates the distribution is hence also a function,

$$\mathcal{G} : \Lambda \rightarrow \Lambda.$$

\mathcal{G} can be decomposed into functions describing the actual genetic operations, i.e., selection, mutation and recombination. This allows for a number of extensions of the model, in order for a more exact characterization of the effects of genetic operations, such as mixing based on an orthogonal basis transform. In the case of finite populations, the transition probability from one population to another can be modeled using Markov chains. Due to its modularity, the framework is not limited to modeling canonical GAs. For example, its extension to time-varying or other search algorithms such as simulated annealing have been demonstrated. The model proved to be suited for the characterization of fixed points (steady states) and metastable states in genetic algorithms. The model has been employed by a number of researchers to analyze different properties of the genetic algorithms, such as the exact schema theorems derived by Vose (1999, chap. 19) and Wright (1999).

4.10.7 Statistical Mechanics Models

In the dynamics system approach proposed by Shapiro and Prügel-Bennett (Shapiro et al., 1994; Prügel-Bennett and Shapiro, 1997), the microscopic behavior of evolutionary algorithms is inferred by assessing the system on a macroscopic level. This approach is very similar to statistical mechanics and thus uses techniques adopted from that field. In contrast to statistical mechanics, however, where the macroscopic quantity is normally the energetic level of a system, here the fitness distribution within the population is taken as the observable factor. Specifically, the goal is to predict the fitness distribution in a subsequent population as a result of the current population and the application of genetic operators. This is done by estimating the change of cumulants under the studied operations. Similarly to moments, cumulants describe the macroscopic properties of a distribution. For example, the first cumulant corresponds to the mean, the second one to the variance of the distribution. Higher order cumulants are related to higher order properties of the distribution: the third cumulant characterizes the (a)symmetry of the distribution, and the fourth one its slope. A number of interesting observations resulted from the application of this framework to genetic algorithms. For instance, selection (the authors used Boltzmann selection) lead to an increase in asymmetry of the fitness distribution, which the authors argued should be compensated by the mixing operators. The inference of the impact of mixing operators, specifically mutation and crossover, is strongly problem-specific. The authors have illustrated the application of their model framework to a number of simplistic test problems. A generalization of their findings to more complex or even arbitrary problems is far from being straightforward.

4.11 Multi-objective Genetic Algorithms

Many optimization problems, especially in technical or other “real-world” applications, involve not only one but a number of different objectives. An example could be the maximization of a performance measure while at the same time resources are to be minimized. Constrained optimization, when realized as a joint optimization of the objective and a constraint satisfaction task, can actually be considered an example of a bi-objective problem (Section 4.4). There, like it is typical in multi-objective problems, the figures of merit are incommensurable, even contradictory. Formally, for unconstrained optimization tasks, we can, WLOG, state the problem as a minimization problem of a function,

$$\mu : X \rightarrow Y,$$

where X and Y are vector spaces, for example, \mathbb{R}^n and \mathbb{R}^m , respectively. The search for a minimum (element) depends on the definition of a partial order on Y .

In numerical optimization, a common approach is to reformulate the multi-objective optimization task as a scalar problem. A typical scalarization technique is the combination of the individual objectives into

a *weighted sum*. The penalty approach for constrained optimization is an example for this procedure. As shown there, the determination of adequate weights is critical for this scalarization approach. The first role of weighting factors is to impose a scaling on the sub-objectives to ensure that all fitness functions yield values of comparable magnitudes. This normalization step is performed in order to prevent an unintentional bias of the search. The second criterion for the selection of coefficients is the actual weighting of the objectives. That is, criteria deemed more important or difficult can be assigned higher weights. Formally, assuming

$$\boldsymbol{\mu} : X \rightarrow \mathbb{R}^m,$$

and given a weight parameter vector $\mathbf{w} \in \mathbb{R}^m$, we can state the (unconstrained) weighted sum replacement problem as:

$$\mu'_{\mathbf{w}}(\mathbf{x}) := \boldsymbol{\mu}^T(\mathbf{x})\mathbf{w} = \sum_{i=1}^m \mu_i(\mathbf{x})w_i. \quad (4.11)$$

Normally, the weights are chosen such that

$$\sum_{i=1}^m w_i = 1.$$

One of the fundamental advantages of this method is that it can be applied to virtually any (numerical) optimization routine without any further adaptation except for the weight assignment. It is hence also applicable to evolutionary algorithms, where it is consequently used routinely. The major drawback, however, is that by seeking an optimum of the replacement fitness function (the weighted sum) it does not account for the multi-objective nature of the problem. For example, compromises between conflicting objectives are assumed to be implicitly solved by the specific balance that is given through the a priori weight parameter settings. The exact properties of the so-obtained solution is not foreseeable. It is well possible that any other solution not reachable through the weights exhibits a well-suited solution.

A different approach to multi-objective problems is hence to employ an order on the search space that does account for the objectives and their relation to each other. One of the well-established order relations that have been applied for this purpose is the Pareto efficiency.³⁹ On a real vector space \mathbb{R}^n and, without loss of generality, assuming a maximization problem,⁴⁰ we can define the order relation as follows: A vector \mathbf{x} dominates another vector \mathbf{y} if and only if each component of \mathbf{x} is greater than or equal to the corresponding component in \mathbf{y} and there is at least one component of \mathbf{x} that is strictly greater than the respective one of \mathbf{y} (see Figure 4.14(b) for an illustration). Formally,

Definition 4.2 (Pareto Dominance Relation) *Given two vectors \mathbf{x} and $\mathbf{y} \in \mathbb{R}^n$. \mathbf{x} dominates \mathbf{y} ($\mathbf{x} \succ \mathbf{y}$) : \iff*

$$(\forall i = 1, \dots, n)(x_i \geq y_i) \wedge (\exists j \in \{1, \dots, n\})(x_j > y_j)$$

The dominance relation is a *strict partial order* relation, since it is

- *irreflexive*: $\neg(\mathbf{x} \succ \mathbf{x})$,
- *asymmetric*: $(\mathbf{x} \succ \mathbf{y}) \implies \neg(\mathbf{y} \succ \mathbf{x})$ and
- *transitive*: $(\mathbf{x} \succ \mathbf{y} \succ \mathbf{z}) \implies (\mathbf{x} \succ \mathbf{z})$

for all $\mathbf{x}, \mathbf{y}, \mathbf{z} \in \mathbb{R}^n$. By Definition 4.2, a Pareto-optimal element of a set can be defined as an element that is not dominated by any other element of the same set, or

³⁹The relation is transferred from an “efficiency” definition that the Irish micro-economist Edgeworth purported to explain the impact of allocation of goods on individuals’ need satisfaction. This “conjecture” was later extended by the Italian economist Pareto, and is in essence still a widely used formalism in neoclassical microeconomics.

⁴⁰Actually, for each component, the ordering direction (ascending or descending) can be defined individually, allowing for a mix of minimization and maximization objectives (Deb, 2001, p. 28).




Figure 4.14: Pareto optimality (adopted from Seifert, 2006b, p. 14), assuming a bi-objective maximization problem with two arguments (a) decision space with the gray-shaded feasible region and a thick blue line denoting the Pareto-optimal solutions. (b) Corresponding sets in the objective space: The thick blue line is called Pareto front. The example vectors **a** and **b** dominate solution **c**. **a** and **b** are incomparable with respect to the dominance relation. However, **a** epsilon-dominates **b**. The definition for the other vectors is as follows (Miettinen, 1999, pp. 15): The ideal vector **i** is comprised of the individually optimized objectives; if feasible, the set of Pareto-optimal solutions consists of only this solution. The utopian vector **u** is defined as an infeasible vector, slightly better than the ideal vector. The nadir vector **n** consists of the objective values that constitute the lower bound (assuming maximization problems) of the Pareto-optimal set.

Definition 4.3 (Pareto optimality) Given a set $S \subset \mathbb{R}^n$ and (S, \succ) , $\mathbf{x} \in S$ is a Pareto-optimal element with respect to S : \iff

$$(\neg \exists \mathbf{y} \in X)(\mathbf{y} \succ \mathbf{x})$$

In a non-empty set, there is at least one element for which this condition holds. Generally, a number of elements in a set are non-dominated, forming a Pareto-optimal set:

Definition 4.4 (Pareto-optimal set)

$$S' := \{\mathbf{x} \in S \mid (\neg \exists \mathbf{y} \in X)(\mathbf{y} \succ \mathbf{x})\}. \quad (4.12)$$

Similarly to the (strict) Pareto dominance, a weak dominance can be defined as follows:

Definition 4.5 (Pareto Weak Dominance Relation) Given two vectors \mathbf{x} and $\mathbf{y} \in \mathbb{R}^n$. \mathbf{x} weakly dominates \mathbf{y} ($\mathbf{x} \succeq \mathbf{y}$) : \iff

$$(\forall i = 1, \dots, n)(x_i \geq y_i)$$

Of course we have that,

$$\succ \subseteq \succeq.$$

For example as a quality measure, as will be shown in Section 4.11.5, it is sometimes useful to determine the proximity of points to the Pareto front. We therefore define a relation, which could be termed “almost dominates” in the following way (see Figure 4.14(b) for an illustration):

Definition 4.6 (Pareto Epsilon Dominance Relation) *Given two vectors \mathbf{x} and $\mathbf{y} \in \mathbb{R}^n$. \mathbf{x} epsilon-dominates \mathbf{y} ($\mathbf{x} \succ_{\epsilon} \mathbf{y}; \epsilon \in \mathbb{R}$) : \iff*

$$(\forall i = 1, \dots, n)(x_i + \epsilon \geq y_i) \wedge (\exists j \in \{1, \dots, n\})(x_j + \epsilon > y_j).$$

The concept of a weak Pareto epsilon dominance is analogously defined as,

Definition 4.7 (Pareto Weak Epsilon Dominance Relation) *Given two vectors \mathbf{x} and $\mathbf{y} \in \mathbb{R}^n$. \mathbf{x} dominates \mathbf{y} ($\mathbf{x} \succ_{\epsilon} \mathbf{y}; \epsilon \geq 0$) : \iff*

$$(\forall i = 1, \dots, n)(x_i + \epsilon \geq y_i)$$

Normally ϵ is assumed to be greater than zero. For ϵ equal to zero, the Pareto epsilon dominance and the dominance relation are equivalent, epsilon values smaller than zero can be interpreted as an indicator for the degree of dominance.⁴¹ The concepts of Pareto optimality and Pareto-optimal sets can be adapted to the Pareto epsilon dominance in a straightforward way.

The application of the former definitions to multi-objective search is also straightforward. Given an unconstrained maximization problem (WLOG)

$$\arg \min_{\mathbf{x}} \boldsymbol{\mu}(\mathbf{x}) \tag{4.13}$$

where, as before, $\boldsymbol{\mu} : X \rightarrow Y$. Following the common terminology used in multi-objective optimization, the optimization variables ($\mathbf{x} \in X$), that is, the domain of $\boldsymbol{\mu}$, are typically termed *decision variables* (Figure 4.14(a)). Accordingly, X is called the *decision (variable) space*. The codomain of $\boldsymbol{\mu}$ is called the *objective space* (Figure 4.14(b)).⁴² Constraints can be incorporated in the same fashion as in single-objective optimization formulation, using equality and inequality constraints. The decision space can then be partitioned into a feasible and an infeasible domain. For symmetry reasons, the same distinction is often exercised in the objective space, calling the feasibility image of ($\boldsymbol{\mu}(X) \subset Y$) the feasible objective space (in game theory sometimes also *multi-loss map*) and the remaining codomain the infeasible objective space.

With these definitions, a global Pareto-optimum can be defined as the set of all feasible candidate solutions that are not dominated by any other feasible solution:

Definition 4.8 (Globally Pareto-optimal set) *Given the (feasible) decision space X and the function $\boldsymbol{\mu}$ as before. \mathfrak{P} is the globally Pareto-optimal set with*

$$\mathfrak{P} := \{\mathbf{x} \in X \mid (\neg \exists \hat{\mathbf{x}} \in X)(\boldsymbol{\mu}(\hat{\mathbf{x}}) \succ \boldsymbol{\mu}(\mathbf{x}))\}.$$

The image of the Pareto-optimal set, is often called Pareto front.

Definition 4.9 (Pareto front) *Given the Pareto-optimal set as defined in Definition 4.8, the Pareto front \mathfrak{F} is defined as:*

$$\mathfrak{F} := \boldsymbol{\mu}(\mathfrak{P})$$

⁴¹We shall see later (Section 4.11.5) that this is an important property for the Pareto dominance compliance of quality indicators.

⁴²In some disciplines such as operations research (OR) a line is drawn between measurable *criteria*, abstract *objectives* and attainable *goals*. We will not use these terms in such a sense here but stick to “objectives” or fitness values in the context of EAs (cf. Coello Coello et al., 2007, p.6).

Similarly as in the single-objective case, local optimality can be defined as the set of solutions that are non-dominated within a (small) neighborhood (e.g., see Deb, 1999; Miettinen, 1999):

Definition 4.10 (Locally Pareto-optimal set/front) (1) The set \mathfrak{P}^* is a locally Pareto-optimal set : \Longleftrightarrow

$$(\exists \epsilon > 0) [(\forall \mathbf{x} \in \mathfrak{P}^*)(\forall \hat{\mathbf{x}} \in X) (|\hat{\mathbf{x}} - \mathbf{x}| < \epsilon \rightarrow \neg[\boldsymbol{\mu}(\hat{\mathbf{x}}) \succ \boldsymbol{\mu}(\mathbf{x})])],$$

where $|\cdot|$ denotes a suitable norm.

In order to impose a neighborhood relation also on the image, the local Pareto front, Zitzler (1999) proposed to require a maximum distance of the function values. Hence the necessary condition (in forward direction) of the definition becomes:

$$(\exists \epsilon > 0)(\exists \delta > 0) [(\forall \mathbf{x} \in \mathfrak{P}^*)(\forall \hat{\mathbf{x}} \in X) (|\hat{\mathbf{x}} - \mathbf{x}| < \epsilon \wedge |\boldsymbol{\mu}(\hat{\mathbf{x}}) - \boldsymbol{\mu}(\mathbf{x})| < \delta \rightarrow \neg[\boldsymbol{\mu}(\hat{\mathbf{x}}) \succ \boldsymbol{\mu}(\mathbf{x})])].$$

(2) Given a locally Pareto-optimal set \mathfrak{P}^* , $\boldsymbol{\mu}(\mathfrak{P}^*)$ is called a local Pareto front.

Under certain assumptions, the scalar sufficient and necessary optimality conditions can be directly expanded to the concept of Pareto optimality—and more general vector optimality formulations (for example, see Kuhn and Tucker (1951), Miettinen (1999) Ehrgott (2005) and Jahn (2011) for detailed discussions).

4.11.1 Multi-objective Optimization Approaches

To utilize the previous definition and the so-resulted Pareto formalism for the solution of multi-objective problems, appropriate optimization routines have to be employed.

Generally, the final result of a multi-objective optimization approach is a set of incomparable solutions, for example, the set of Pareto-optimal solutions. But typically a decision maker (DM) (human or automaton) aims at reducing this set to a feasible size, often only a single solution. The criterion by which such a reduction is performed is termed preference. Multi-objective optimization approaches can be categorized according to the way the preference information is incorporated into the search. Miettinen (1999), for example, proposed the following taxonomy:

1. *No-preference* methods—no explicit DM preference is taken into account, but an implicit preference (e.g., using heuristics) is assumed.
2. *A posteriori* methods—decision making is performed subsequent to the actual optimization that yields a set of optimal solutions.
3. *A priori* methods—the DM expresses preferences prior to search, which is then guided accordingly.
4. *Interactive* methods—the DM interactively influence the optimization direction through iterative preference articulation.

A number of common “classical” approaches will be shortly discussed in the following paragraphs.

A Priori Preference Articulation

With these methods, the decision making process is defined prior to the actual optimization process. In other words, the DM determines preferences before starting the optimization.

In *goal programming* techniques, this is done by defining a goal, typically as a set of objectives, which do not necessarily have to be attainable (Charnes and Cooper, 1955; Ignizio, 1976). The scalar replacement problem can hence be formulated as a minimization of the deviation of observed and intended objective values. This allows for a reformulation of the problem as a linear combination of under and over-attainment values, rendering it amenable to linear programming techniques. Weight parameters are often included in order to rank the importance of individual objectives and the preference of over or under-attainment. The main issue with it is the predetermination of viable goals. For example, with little or without any a priori knowledge of the search space, the predefined goals generally do not represent the best possible compromise solution. Like in weighted sum approaches, the objectives need

to be normalized in order to prevent a non-intended bias to strongly influence the search direction. *Lexicographical goal programming* (Charnes and Cooper, 1961) can be considered a sequential goal programming technique in which the objectives and the associated goals are prioritized according to their importance. In each step, starting with the most important goals, one level of goals is taken into account, yielding a set of solutions. These solutions serve as constraints for the subsequent step in which the next level of goals is optimized for. Once a unique solution is found, the algorithm is terminated. This method is thus not suited to cover the entire Pareto front. It also fails to achieve solutions in some non-convex regions. Moreover, the search is highly biased by the priority of goals; that is, changing their order may lead to entirely different final results. Various other variants of goal programming exist, including the formulation as a *min-max* problem and fuzzy (also Chebyshev) goal programming (Flavell, 1976).

A noteworthy variant, partly aiming at the problems associated with goal programming, is called *goal attainment method* (Gembicki and Haimes, 1975). There, both a goal and a weight vector, determining the relative importance of the goals, are provided by the decision maker, and the problem is restated as a minimization of a slack variable that relaxes the feasible region. This method is particularly well suited for non-attainable goals. The weight vector can then be thought of as the vector pointing toward the feasible (attainable) region, originating from the goal vector. The feasible region is contracted, eventually converging to the final solution, which is given by the intersection of the goal and weight vector combination and the feasible region boundary. Additional (inequality) constraints can be included in a straightforward manner. In contrast to the weighted sum approach, this method allows for non-dominated points in non-convex regions to be found. Moreover, by applying an iterative scheme with varying weight parameters, the method is suited to yield multiple samples of the Pareto front, turning it into a posterior method.

The scalar replacement problem posed by *utility function methods* (Keeney, 1969; Keelin, III, 1976) is the maximization of a utility function that maps the feasible objective space onto a scalar space, for example, of real values:

$$U : \mu(X) \rightarrow \mathbb{R}.$$

In this sense, it can be considered a generalization of the weighted sum or the weighted metric approach, as pointed out by Rosenthal (1985). The function U , however, is assumed more generally as a smooth, differentiable and concave function, yielding a utility measure (e.g., of a decision maker) associated with a point in the objective space. In order for the utility function to express preference, the following definition is imposed:

$$\begin{aligned} \text{the decision maker prefers the outcome of } \mathbf{x}_1 \text{ over that of } \mathbf{x}_2 : &\iff \\ U(\mu_1) &> U(\mu_2), \end{aligned}$$

with $\mu_i := \mu(\mathbf{x}_i)$. The optimization problem then reduces to finding a maximizer of $U(\mu)$, of course, subject to feasibility of μ . The actual problem is the determination of a suited such utility function, which reflects the preference structure of a decision maker. In microeconomics, where this approach is adopted from, the problem is posed as an optimal allocation of monetary resources on multiple different commodities. Each commodity has a price and is available in continuous quantities. This problem can be stated as a maximization of U subject to the sum of prices times quantities being equal to the available fund. The equivalent Lagrange multiplier system under the first order necessary optimality conditions gives the following multipliers for any pair of components of a maximizer μ^* :

$$\lambda_{i,j} = \frac{\frac{\partial U(\mu^*)}{\partial \mu_i^*}}{p_i} = \frac{\frac{\partial U(\mu^*)}{\partial \mu_j^*}}{p_j}, \quad (4.14)$$

where p_i, p_j is the price associated with the respective commodity. Rearranging (4.14) gives:

$$\lambda_{i,j} = \frac{\frac{\partial U(\boldsymbol{\mu}^*)}{\partial \mu_i^*}}{\frac{\partial U(\boldsymbol{\mu}^*)}{\partial \mu_j^*}} = \frac{p_i}{p_j},$$

Thus with very limited notion of U , the optimal trade-off between any two commodities can be characterized. In economics, the individual partial derivatives are called *marginal utility* and its ratio forms the *marginal rate of substitution*. Inversely, if decisions leading to a specific allocation are assumed optimal with respect to the preferences of a decision maker, then an approximation of U can be derived by an assessment of these decision points.⁴³ The utility function can hence be synthesized by employing generic parameterized forms that are specialized through the assessments. In addition, interactive methods, in which within intervals, the decision makers have to express their preference (for example, Geoffrion et al., 1972). In so doing, this method employs implicit, local approximations of U rather than assuming the validity of a global utility function. Moreover, Rosenthal (1985) showed that other multi-objective approaches, such as aggregation methods or goal programming, can be regarded as implicitly making use of utility functions. He argues that the utility functions implied there are comparatively simple and thus may fail to capture non-linear preference structures, such as optimal points on non-convex optimality regions.

A Priori and A Posterior Preference Articulations

For some approaches, preferences can be expressed either prior or subsequent to the actual optimization process. The *weighted sum* technique, for example, is traditionally employed with a predefined weight configuration. Applying an iterative scheme with variable weight settings, a Pareto front approximation can be obtained, rendering it an a posteriori approach.⁴⁴ There are, however, several problems associated with this procedure. The main drawback is that it fails to find any solution in non-convex regions of the Pareto front. This can be demonstrated following the illustrative justification (for a bi-objective case) by Das and Dennis (1997). There, an equivalent problem using a rotated coordinate system is considered. In the convex case (see Figure 4.15a, where a maximization problem is assumed), each weight configuration corresponds to a single maximum. By changing the weight parameters, which corresponds to altering the rotation angle θ of the coordinate system, any solution on the Pareto front can be obtained. In the non-convex case (Figure 4.15b), points in concave regions such as R do not maximize the equivalence problem. By shifting the dotted line, which corresponds to a specific weight configuration, a better performing solution is achieved. In this example, two solutions, Q and P , yield an equal merit, revealing yet another issue of the weighted sum approach: Unique solutions under Pareto optimality may become ambiguous when compared using weighted sums. Finally, when approximating the Pareto front, it is highly desirable to obtain an uniform spacing of the sample points. Generally, because of the non-linear relationship between weights and the corresponding pay-off, this cannot be guaranteed using weighted sums.

Reference point methods (Wierzbicki, 1980) and *compromise programming* approaches (Yu, 1973) are closely related to the previously introduced *goal programming* approach (Romero et al., 1998); thus rendering it a priori preference methods. The optimization task is reformulated as a minimization problem of the deviation of the observed fitness from an ideal solution. Commonly, this ideal solution is an infeasible point in the objective space such as the ideal or utopian vector (\mathbf{i} and \mathbf{u} in Figure 4.14(b)). The optimization procedure can be described as finding the attainable solutions that is closest to this ideal point. This method strongly depends on the selection of weights and on the distance metric. The

⁴³Remarkably enough, these considerations constitute the basis of a significant portion of modern economics, in which specific outcomes and the relation of these outcomes are taken as an explanation of the underlying laws. Nonetheless, it is this area of research that has contributed a great deal of methods and theories to the field of multi-variate and multi-objective optimization, which in turn is of great *utility* for a variety of disciplines.

⁴⁴Any solution to the weighted-sum replacement problem are Pareto-optimal, provided that all weight parameters are positive (Deb, 2001; Miettinen, 1999; Zadeh, 1963).

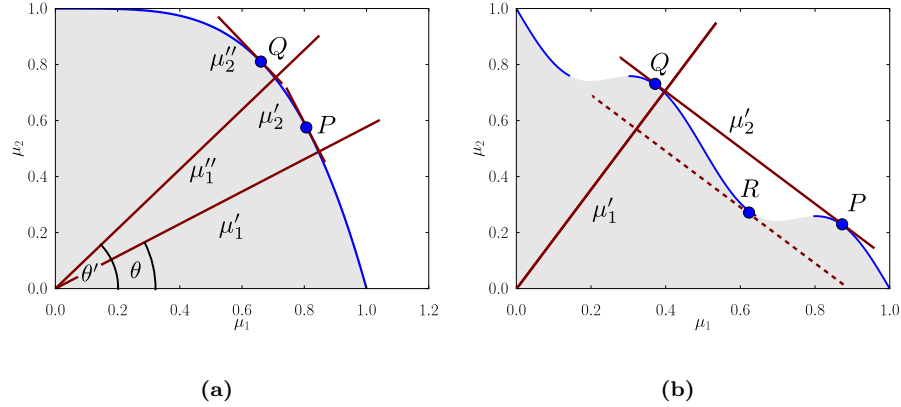


Figure 4.15: Weighted sum maximizers for a bi-objective optimization problem. The objectives are denoted by μ_1 and μ_2 . Two different weight configurations that correspond to (counterclockwise) coordinate system rotations by angles θ and θ' , respectively. The so-obtained coordinate systems are marked (μ'_1, μ'_2) and (μ''_1, μ''_2) . (a) In the convex case, the equivalent maximization problem yields one unique solution per weight setting. Also, for each point on the Pareto front, there exists a weight configuration. (b) In the non-convex case, however, Pareto-optimal solutions in concave regions do not maximize the weighted sum and are hence abandoned. Moreover, one weight configuration may lead to multiple maxima (example after Das and Dennis, 1997).

latter becomes even more critical when the method is used as a posterior, sequential technique for the approximation of the Pareto front (see, for example, Kaliszewski, 1987). There, similar problems as with weighted sums occur; that is, a deficiency to find all Pareto-optimal points and a non-uniform sampling of the front. Related iterative approaches were proposed in which different reference points are adaptively generated in order to achieve a higher coverage and uniformity of the sampled Pareto front. Especially the *reference point methods* and the *light-beam search* variants are often applied as interactive methods, in which the DM can iteratively refine reference objectives (Wierzbicki, 1980; Jaskiewicz and Słowiński, 1999).

A Posterior Preference Articulation

The following methods can be said to decouple the optimization and the decision making process, which can be performed in a post-optimization procedure, for example, using visualization-aided (Deb, 2001, pp. 316–320), (Miettinen, 1999, chap. 3) manual selection methods or automated selection procedures—for example, based on an a priori method Section 4.11.1 (Deb, 2001, pp. 390–392).

In the *epsilon-constraints method*, the replacement problem is given as a constrained optimization problem (Marglin, 1967; Haimes, 1971). This is done by keeping one of the objectives as the function to be optimized and introducing inequality constraints for the remaining objectives. In order to sample the Pareto front, the problem is solved for different epsilon values. In each step, in which the epsilon parameter vector is fixed, the intersection between the Pareto front and the feasibility hyperplane is sought, yielding a unique Pareto-optimal solution. This procedure is repeated, for example, until an appropriate approximation of the Pareto front is obtained. This method is suited to potentially find any Pareto-optimal solution, regardless of whether it lies in a convex or concave region. The main drawback is that it depends on a user-provided predetermination of the epsilon parameters, thus requiring some

a priori knowledge of the search space. Adaptive epsilon determination methods, aiming at an ideal Pareto front sampling have been proposed, however (see, for example, Eichfelder, 2006).⁴⁵

Another approach that iteratively generates Pareto-optimal points was developed by Das (1998) and coined *normal boundary intersection* method. The main idea can be broadly summarized as follows: Given any Pareto-optimal points—as anchor points, the authors propose to use minimizers that are global minimizers of exactly one of the objectives—and the convex hull of these points: In many situations, the intersection of a line segment, due to the normal of the convex hull and any point inside, and the Pareto front forms another Pareto-optimal point. In fact, the only situation in which the former does not hold is when the Pareto front is folded so that a part of the front becomes shadowed by a sub-optimal section. The optimization can then be reformulated as a search for these intersections under different convex combinations that are given by different weight parameters. The authors show that with an appropriate ordering, the intersection-finding subproblems can be solved in an efficient, hierarchical fashion. Moreover, in contrast to the weighted sum approach, a uniform sampling of the Pareto front can be achieved.

A similar approach, called *normal constraint boundary* method was presented by Messac et al. (2003). There, after generating anchor solutions like in the normal boundary intersection method, the objective space is normalized in order to achieve uniformly distributed Pareto solutions. The hyperplane spanned by the anchor solution is equidistantly partitioned. The intersection between the normal arising from each of the partitions and the Pareto front is located, and the so-generated points are evaluated according to their Pareto optimality. They may fall into three categories, globally Pareto-optimal, locally Pareto-optimal or non-Pareto-optimal. To retain only globally Pareto-optimal solutions, the authors devise an according filtering technique, which can also be applied to other related techniques such as the normal boundary intersection method. Recently, an application of a genetic algorithms to generate the non-dominated solutions has been proposed (Martínez et al., 2009). The authors claim that this approach renders the use of any filtering technique obsolete.

4.11.2 Evolutionary Multi-objective Approaches

Many of the techniques described above are in general amenable to using an evolutionary algorithms as their underlying search technique. The weighted sum approach, applied as an a priori method, is indeed still the most common approach to multi-objective optimization, in spite of its deficiencies. Besides the fact that this form of aggregation of different traits is often considered “natural,” the following rationales can be associated with its utilization:

- (1) It does not involve any adaptation of the evolutionary algorithm.
- (2) In many situations the feasible objective space is fully convex, such that all solutions are principally reachable.
- (3) Decision making, that is, selecting the final solution from the set of non-dominated solutions, is not always considered a blessing but rather an extra burden since it may interrupt an otherwise automated flow, for example, by requiring user-interaction or a tailored post-optimization procedure.
- (4) When interested in a unique solution, the sampling of the Pareto front does not matter, albeit the non-linear relationship between weights and the position of the solution on the front may.
- (5) The computation time associated with a posteriori methods is in general higher compared with a priori ones.

Often, however, the limitations associated with these methods are critical, prompting for a dedicated multi-objective approach instead. Genetic and other evolutionary algorithms, being population-based routines, can be considered well-suited for the tasks associated with the mutual optimization of different objectives and the two critical goals associated with it: positioning the final solutions on the Pareto front or at least in close proximity and achieving a high coverage of the entire front. In contrast to most other optimizers, these goals can potentially be solved implicitly and in parallel. It also introduces

⁴⁵The method developed there is based on a scalarization technique proposed by Pascoletti and Serafini (1984) in which an optimality criterion based on the partial ordering induced by a convex cone, a more general definition than Pareto optimality, is employed and of which the epsilon-constraints method can be shown to be special case (also see Yu, 1974).

a number of additional issues, however, that have to be taken into account. As we have seen earlier, genetic algorithms face a conflict often described as trade-off between exploration and exploitation. This is so with single-objective optimization, where this balance only impacts the distribution of candidate solutions in the decision space, and can hence be expected to become more critical in the multi-objective case, where the distribution of solutions in the objective space has to be considered, too. Consequently, niching techniques (as explained in Section 4.9) play a paramount role in multi-objective evolutionary algorithms (MOEAs).

A number of straightforward approaches, incorporating methods described above but exploiting the population nature of GAs have been suggested. As shown above, using sequential schemes with different parameter configurations, the weighted sum method can be used as an a posteriori method. Different approaches to incorporate a similar scheme, exploiting the parallelism of GAs have been proposed (cf. Deb, 2001, pp. 189–200). Hajela and Lin (1992), for example, proposed an a posteriori weighted-sum-based GA, in which the chromosome of each solution is supplemented with a different weight vector. In order to maintain diverse non-dominated solutions, a niching technique, restricted to the part of the chromosome that represents the weight configurations, is employed. Alternatively, the authors propose different sub-population approaches in which each population is evaluated using different weight settings. As another alternative to niching or sub-population approaches, Murata and Ishibuchi (1995) propose randomized weight vectors. Thus each evaluation is performed with a different, random weight configuration, avoiding any premature weight bias. In addition, the authors employ an elitist selection scheme in which an external archive containing non-dominated solutions is maintained supplementing the main population. We will see later that many of the state-of-the-art approaches employ a similar technique to incorporate elitism. Despite its efficacy in the context of various, especially combinatorial, problems, these GAs face the same problems as other weighted-sum approaches. Most importantly, they fail to find solutions in non-convex regions of the Pareto front and they are prone to incompletely covering the Pareto front.

Also the epsilon-constrained replacement problem is principally amenable to being solved using GAs. Ranji Ranjithan et al. (2001) propose a sequential scheme in which the individual Pareto-optimal solutions are obtained through evolutionary algorithm runs. Instead of actually starting off with a randomly initialized population, however, the restarted EA instances re-use the population resulted from the previous run. The idea behind this procedure is to exploit potential decision space similarities of adjacent Pareto-optimal solutions. Yet, in order to increase the exploratory power of the early phase of each EA run, the authors employ initially high but time-decreasing mutation rates. The extreme points and the step-width of the epsilon constraints are user-provided. As also noted by Landa Becerra and Coello Coello (2006), no information about the underlying constraint handling technique is provided. Landa Becerra and Coello Coello (2006) apply their (see Section 4.4) cultured differential evolution method to sequentially obtain Pareto-optimal solutions. Their approach, like all epsilon-constraints methods, depends on an a priori determination of the epsilon constraints and hence on their lower and upper bounds, for which the authors provide an automatic approach based on the ideal and the nadir vectors. The latter can be approximated in a straightforward manner only for bi-objective problems—a meta-heuristic for higher dimensions was proposed by Laumanns et al. (2006). Although potentially suited to solve also hard multi-objective problems, this approach may become computationally expensive, since it requires the sequential execution of multiple evolutionary algorithms. Moreover, it strongly depends on an effective and efficient constraint handling mechanism, which can easily turn into a cyclic problem, for example, when regarding constraint handling as a bi-objective problem consisting of the actual optimization and the constraint satisfaction.

One of the early approaches solving multiple optimization criteria with a standard GA is the vector-evaluated genetic algorithm (VEGA) developed by Schaffer (1985). It incorporates multiple populations, each of which represents a different object. That is, given a population size N , and m objectives, each objective is accounted for by selecting a fraction of N/m of the population according to the fitness associated with it. This procedure does not exercise a strict partitioning on the population as a candidate solution may of course exhibit above-average fitness values for several objectives. The basic idea is that offspring individuals—albeit their parents were selected from the extremes, that is, high fitness values in one of the objectives—eventually constitute compromise solutions. It was found, however, that without

extra effort, a diverse sampling of the Pareto front could not be achieved. Schaffer tried to foster the generation of compromise solutions by mating restrictions, in which mating of individuals fit in terms of different objectives was encouraged. Additionally, locally ideal compromise solutions (i.e., locally Pareto-optimal solutions) were rewarded in order to increase their selection probabilities. With these additions, the method proved effective for simple problems with convex objective spaces.⁴⁶ It is also decently efficient, if a good coverage of the Pareto front is of minor concern (Zitzler et al., 2000). To date, the main application of VEGA is for constraint handling (Coello Coello et al., 2007).

The majority of state-of-the-art evolutionary algorithms use a ranking mechanism that assigns a non-dominated rank to each candidate solution in a population. The so obtained ranks can then be readily used in selection. That is, the number of trials an individual receives is proportional to its non-dominated rank. *Non-dominated ranking* as schematically demonstrated in Algorithm 4.4 was proposed by Goldberg (1989), and he describes it as technique in which the non-dominated solutions are peeled off from the population rank by rank. In other words, first, all non-dominated solutions are assigned a rank of one and removed from the original population. The same procedure is repeated for the remaining individuals, giving them a rank of two. The routine is repeated until the original population is empty and, consequently, all solutions have been ranked. Non-dominated-ranking can

Algorithm 4.4: Basic non-dominated sorting algorithm

Data: P // population

```

1  $r \leftarrow 0$ 
2 while  $|P| > 0$  do
3    $N \leftarrow$  non dominated solutions in  $P$ 
4   foreach  $A \in P$  do
5      $R(A) \leftarrow r$ 
6    $P' \leftarrow P' \cup N$ 
7    $P \leftarrow P \setminus N$ 
8    $r \leftarrow r + 1$ 
```

be included in virtually any evolutionary algorithm by incorporating the scheme above, using a suited selection operator—such as tournament selection or rank-based selection schemes—that allows for the non-dominated rank to be used as selection criterion. Without further modifications, such a procedure leads to the following problems: (1) There is no differentiation between solutions with the same rank, yielding an arbitrary sampling of the Pareto front. (2) The average selection pressure diminishes with the number of objectives since the number of incommensurable solutions increases.

Fonseca and Fleming (1993) were the first to devise a GA [called multi-objective genetic algorithm (MOGA)] using a similar non-dominated sorting scheme as proposed by Goldberg. In contrast, however, their sorting algorithm distinguishes solutions not only by the level of non-domination but also by the number of individuals by which a solution is dominated:

$$R(A) = 1 + |\{B \mid B \succ A\}|.$$

Like in Goldberg's approach, non-dominated solutions receive a rank of one. But the rank distribution of dominated candidates is such that individuals that are dominated by fewer solutions are slightly preferred. To additionally maintain diversity of solutions with the same rank and thus to eventually ensure a broader coverage of the Pareto front approximation, the authors employ a niching technique to maintain the diversity of non-dominated solutions. They propose a fitness sharing scheme (Section 4.9.1), where the distance between solutions of the same rank is evaluated in the objective instead of the decision space. As the selection operator, a stochastic universal sampling scheme was incorporated, and a linear interpolation transforming the ranks in fitness values is used. The main disadvantage of this

⁴⁶As noted by Richardson et al. (1989), the method is principally no different from weighted scoring approaches in that it linearly combines the objectives vectors, albeit with implicit weight configurations, given by the current population.

approach is the fact that the fitness assignment does not explicitly prefer solutions with lower ranks (Deb, 2001, p. 206). For example, low-ranked solutions (i.e., solutions near local optima) in crowded regions may exhibit a lower fitness than highly dominated one in sparse areas. This may introduce a non-intended bias and impair the convergence to the Pareto front. The use of a different crowding strategy, closely resembling restricted tournament selection (see Section 4.9.2), was later proposed by Andersson and Wallace (2002). Their multi-objective struggle GA is specifically aimed at maintaining diverse local Pareto fronts.

Fitness-proportionate selection was employed in the former approach in order for fitness sharing to be applicable. To directly select solutions based on the ranks that result from non-dominated sorting, the niched-Pareto genetic algorithm (NPGA) by Horn et al. (1994) incorporates tournament selection in conjunction with a dynamic sharing mechanism similar to the one proposed by Oei et al. (1991)—explained in Section 4.9.1. Specifically, a binary tournament is performed in which, like usual, the two competitors are randomly picked from the current population. In addition, a sub-population, consisting of (t_{dom}) individuals, is randomly drawn from the current population, and the two tournament individuals are compared with them concerning their dominance. Three situations may result: (1) One candidate of the tournament is not dominated with respect to the reference sub-population, the other one is. In that case, the former is selected for mating. (2) Neither candidate is dominated. (3) Both candidates are dominated. In the two latter cases, sharing is performed to determine the winner of the tournament. This is done by computing the niche count for each candidate in the offspring population, which is being built. The solution with the lower niche count wins the tournament. Like in the MOGA case, the niching distance is evaluated in the objective space. This procedure dynamically prefers solutions whose objective values are underrepresented in the target population. Conventional crossover operations are performed between the winners of two tournaments, optionally followed by mutation. For the first two tournaments, the target population is hence empty. Consequently, the niching-based tournament, if it applies [situations (2) or (3)], cannot be performed for the first two parents, in which case a random selection is used instead. NPGA has been shown to perform well, especially on problems with many objectives. But it also strongly depends on the tournament comparison windows size (t_{dom}) and on the sharing threshold σ (Deb, 2001; Zitzler et al., 2000).

As discussed in Section 4.5, elitist selection has been found to improve the convergence behavior of single-objective EAs (De Jong, 1975) and is also a sufficient condition for their convergence (Rudolph, 1996); also see Section 4.10.5. It thus stands to reason to incorporate similar concepts into MOEAs. Unfortunately, a concept similarly straightforward as in the single objective case cannot be devised. For example, because of the optimality criterion, the number of equally well paying off individuals in a populations cannot be predetermined, and so an adequate generational gap cannot (for a discussion, see Deb, 2001, pp. 239–240). An early attempt to account for this difficulty in incorporating elitism was made by Osyczka and Kundu (1995). Like many state-of-the-art approaches, discussed later, they propose the use of a secondary population, often termed archive. The archive consists of solutions not dominated by any solutions of the population up to the current generation. That is, solutions are transferred into the archive, but there is no direct back-propagation of elite solutions into the primary population, which is in contrast to conventional elitist selection schemes. Instead the archive serves as a reference for the fitness assignment that is based on the relative distance of solutions from the local Pareto front. Specifically, a new solution that is not dominated by any archived candidate is rewarded proportionally to the distance it exhibits to the closest archive member. It is added to the archive. Additionally, if it itself dominates any solutions in the archive, those candidates are removed. Contrariwise, if the new solution is dominated by any archive member, it is penalized with the distance to the closest member. Obviously, the purpose of this mechanism is to assign greater fitness values to non-dominated solutions. At the same time, it is aimed at increasing the diversity of (locally) Pareto-optimal candidates, since it assigns a greater fitness to those optimal solutions that are far from previous non-dominated individuals. Deb (2001, pp. 258–259) shows that under certain conditions this can lead to an opposite effect, emphasizing crowded regions. Deb (ibid.) also argued that the approach should be supplemented with a size limitation of the archive, which otherwise might continuously grow and lead to performance issues.

Another class of approaches that make use of an external archive are Pareto-archived evolution strategy (PAES) and Pareto envelope-based selection algorithm (PESA). The former is an evolution strategy (Knowles and Corne, 2000), and $(1 + 1)$ and $(\mu + \lambda)$ replacement schemes have been demonstrated (cf. Section 4.1.1). In contrast to the former approach, the external archive is size-limited. Thus for new individuals eligible to enter the archive, that is, those that are not dominated by any archive members, a replacement strategy is to be defined; for example, as the following two-fold fashion: (1) If the solution dominates one or more archive members, this solution enters the archive, and the dominated members are removed. (2) Otherwise, if the archive is not yet completely filled, the solution enters the archive. If, on the other hand, the maximum archive size would be exceeded with the new solution, a density evaluation is performed. If the candidate fills a highly dense slot with respect to the objective space, it is rejected. But if the solution would be situated in a sparse region in the archive, another member of the archive in a more crowded region is removed. The candidate enters the archive instead. Diversity preservation is achieved through a clustering technique (cf. Section 4.9.1): The density estimation is performed by subdividing the objective space into hypercubes. Each of these hypercubes is assigned a squeeze value to, reflecting the number of solutions it is inhabited by. The higher this number, the higher the probability of solutions of this cube to get erased from the archive. The same density estimator is used for the replacement in the primary population. In the $(1 + 1)$ strategy, if parent and offspring solutions do not dominate each other the one located in a sparser region with respect to the archive is preferred. Otherwise the dominating candidate becomes the parent solution of the next generation. In $(\mu + \lambda)$ strategies, the same concept is used to assign fitness values to individuals. After the λ children have been generated (in the original work through mutation only) a domination marker is set for all of the parent and offspring solutions. This marker can take three states: (1) solution is dominated by any one solution in the archive, (2) solution is neither dominated nor dominates or (3) solution dominates one or more solutions of the archive. Fitness values are assigned corresponding to the markers. That is, regardless of the position in the objective space, dominating solutions will always receive the higher fitness values than non-dominated solutions, which in turn exhibit higher payoffs than dominated ones. Solutions with the same dominance score are evaluated in terms of the location in the grid; solutions in sparse regions are rewarded whereas those in dense regions are penalized. Similarly to the normalization step in fitness sharing, the authors propose the use of an adaptive grid that takes the range of the current objective function values into account. PAES has proved successful on a number of simple test problems.

The same strategy has been applied to a genetic algorithm (Corne et al., 2000). Called PESA, it makes use of binary tournament selection, where the fitness is assigned in the same manner as in PAES. According to the tests conducted by the authors, it outperformed earlier versions of two other state-of-the-art approaches, non-dominated sorting genetic algorithm (NSGA) and strength Pareto evolutionary algorithm (SPEA), whose successors will be discussed in the following sub-sections. Later, a revised version of PESA, PESA-II (Corne et al., 2001), was developed, in which tournament selection is not performed on individuals but on the hypercubes. Specifically, hypercubes receive fitness values based on their location with respect to the local Pareto front and on the squeeze factor—the number of solutions that it houses. After selecting a hypercube according to its fitness and using any standard selection scheme, such as tournament selection, any solution located in this hypercube is randomly chosen. The rationale behind this approach—coined *region-based selection*—is that with individual-based tournament selection the probability of choosing a candidate from sparse regions becomes dramatically low, thus thwarting the diversity-preserving effort at which the clustering technique is aimed. Competitive results of this approach have been reported by the authors who compared it with earlier versions of SPEA and the authors' predecessor implementations, PAES and PESA.

Similarly as in the single-objective GA case (see Section 4.10.3), the building block genesis can be explicitly accounted for in MOEAs. Van Veldhuizen and Lamont (2000) have proposed an according algorithm that can be characterized as a straightforward transformation of the messy genetic algorithm approach to the scope of multiple objectives—called multi-objective messy GA (MOMGA). A detailed discussion and assessment of the MOMGA and its extensions was conducted by Zydallis (2003).

The next two sub-sections are dedicated to introducing two popular state-of-the-art algorithms, NSGA-II and SPEA 2, which have been used in this work.

4.11.3 Non-dominated Sorting Genetic Algorithm

The non-dominated sorting genetic algorithm is one of the earliest approaches using a ranking procedure based on the Pareto dominance relation, which has been first proposed by Goldberg (1989). Like many related approaches that were developed in the beginning of the 1990s, the original NSGA (Srinivas and Deb, 1994) is concerned with maintaining a high degree of objective-space diversity but exhibits a relatively slow convergence behavior, which can be attributed to the lack of elitism, that is, Pareto-optimal solutions are not guaranteed to survive genetic operations. Moreover, the run-time complexity is in the order of $O(MN^3)$ — M is the number of objectives and N is the population size—caused by the high number of mutual dominance relation assessments required to rank the population. It is summarized in Algorithm 4.5. This ranking routine can be coarsely described as follows: A temporary

Algorithm 4.5: NSGA population ranking

Data: P // population

```

1  $r \leftarrow 0$ 
2 while  $|P| > 0$  do
3   foreach  $\mathbf{p} \in P$  do
4      $d_{\mathbf{p}} \leftarrow \text{False}$ 
5     foreach  $\mathbf{q} \in P \setminus \{\mathbf{p}\}$  do
6       if  $\mathbf{q} \succ \mathbf{p}$  then
7          $d_{\mathbf{p}} \leftarrow \text{True}$ 
8         break
9     if  $d_{\mathbf{p}} = \text{False}$  then
10       $R(\mathbf{p}) \leftarrow r$ 
11       $P \leftarrow P \setminus \{\mathbf{p}\}$ 
12    $r \leftarrow r + 1$ 

```

copy of the current population is created. All individuals of this population that are not dominated by any other solution are assigned a rank of zero. These non-dominated individuals are then removed from the temporary population. First-rank solutions are then determined in the same way, albeit on the population without the formerly removed rank-zero solutions.

The run-time complexity of this procedure, which is widely used also in other multi-objective optimization approaches, is of major concern as it drastically affects the overall run-time behavior of the optimizer.⁴⁷ Therefore Deb et al. (2002a) proposed a more time-efficient sorting algorithm that makes use of a hierarchical ranking and that has been integrated into a modified version of NSGA, termed NSGA-II. The algorithm has a run-time complexity of $O(MN^2)$ and can be characterized as the following two-stage procedure:

- (1) For each solution, the number of other solutions that dominate it is determined, and, in addition, a set containing those individuals that it dominates is created. Non-dominated solutions—those with a counter equal to zero—receive the rank zero.
- (2) The rank-one set is constructed by visiting all solutions dominated by rank-zero individuals. At each visit, an individual's counter is decreased by one. A counter value of zero means that the individual is dominated only by rank zero individuals; consequently, the respective individual receives a rank of one. In order to determine individuals with a rank of two, the second step (2) is repeated, now for all solutions dominated by rank-one individuals. Finally, the second step is iteratively called until all individuals have been ranked.

Alternatively, as proposed by Seifert (2006b), the ranking algorithm can be regarded as a topological sorting algorithm on the directed acyclic graph (DAG) spawned by the dominance relation. Jensen

⁴⁷Although in this work and many other real-world optimization problems, in which the merit evaluation involves long-running numerical computations, it can be argued to be of marginal interest.

(2003) proposed a divide-and-conquer algorithm with which the time complexity of the non-dominated sorting routine can be reduced to the order of $O(MN \log^{M-1} N)$, where N is the population size and M is the number of objectives. Fang et al. (2008), however, showed that Jensen's algorithm is not equivalent to the NSGA-II version: While the original algorithm assigns different ranks to duplicate solutions, with Jensen's proposal they are considered to lie on the same front. As an alternative, Fang et al. developed an algorithm also adopting a divide-and-conquer technique but additionally exploiting the hierarchical nature of the dominance relation, using a DAG structure,⁴⁸ which is additionally suited to prevent multiple comparisons for one pair of individuals.

In addition to a fast convergence to the set of Pareto-optimal solutions, one critical goal of multi-objective optimization approaches is a wide coverage and a uniform sampling of the Pareto front. For that purpose, the original NSGA implementation employs the fitness sharing technique, explained in Section 4.9.1. The performance of sharing strongly depends on the distance threshold parameter σ . In addition, sharing leads to a runtime of the order $O(N^2)$ since, in the worst case, it involves comparisons of each individual with all other individuals. In the revised version of NSGA, the sharing function has hence been replaced with a phenotypic⁴⁹ crowding approach (cf. Section 4.9.2), which circumvents both the dependence on user-provided sharing parameters and the quadratic runtime behavior. The crowding approach proposed by Deb et al. (2002a) uses a density estimator, referred to as crowding distance. The density for one candidate is estimated by constructing a hyperrectangle whose dimensions are given by the candidate's distance to the adjacent solutions with respect to one of the objectives. This is illustrated in Figure 4.16. To determine adjacent solutions regarding one objective, the population has to be sorted with respect to that objective, thus involving the same number of sorting steps as there are objectives.




Figure 4.16: Density estimate used in NSGA-II. The density of the solution candidate A_i is determined by constructing the hyperrectangle whose dimensions are given by the adjacent candidates with respect to one of the objectives. The average side length of the hyperrectangle then serves as a density metric.

The so-obtained density measure is used in an ordering relation called *crowding comparison operator* (\prec_n), defined as follows:

$$A \prec B : \iff (R(A) < R(B)) \vee (R(A) = R(B)) \wedge (d(A) < d(B)),$$

⁴⁸More precisely, they employ trees with edge labels denoting the dominance relation of solutions.

⁴⁹The authors argue that it could be replaced with genotypic crowding in straightforward manner.

where $d(\cdot)$ is the density estimator, which is reciprocally proportionate to the dimensions of the hyper-rectangles and $R(\cdot)$ denotes the rank of individuals. In other words, the crowding comparison operator supplements the rank-based comparison with a secondary density-based comparison.

Another extension of NSGA-II is the introduction of elitism. In contrast to related approaches—such as the SPEA, discussed in the next sub-section—it dispenses with an external archive. Instead a similar technique as the $(\mu + \lambda)$ scheme in evolution strategies (cf. Section 4.1.1) is employed (see Algorithm 4.6). After one generation, the intermediate population consists of parents and the children, resulted from crossover and mutation. The intermediate population is then decomposed into its ranks (R_0, R_1, \dots, R_l) . The new population is generated by incorporating as many complete rank sets from the intermediate solution as possible. Since the inclusion of the last-considered rank set may lead to an overfull population, only the fraction that is optimal according to the crowding comparison operator is added. As a result, the new population thus contains the first μ Pareto and crowding-optimal candidates of the intermediate population.

Algorithm 4.6: NSGA population reduction

```

Data:  $P^{(k-1)}$  // parents  $|P^{(k-1)}| = \mu$ 
Data:  $C$  // children  $|C| = \lambda - \mu$ 
1  $P' \leftarrow P \cup C$ 
2  $P^{(k)} = \emptyset$ 
3  $R_0, R_1, \dots, R_l \leftarrow \text{non-dominated-ranking}(P')$ 
4  $i \leftarrow 0$ 
5 while  $|P^{(k)}| < \mu$  do // next population not complete
6   if  $|P^{(k)} \cup R_i| \leq \mu$  then // next rank entirely fits into next population
7      $P^{(k)} \leftarrow P^{(k)} \cup R_i$ 
8   else
9      $R_i \leftarrow \text{non-dominated-crowded-sorting}(R_i)$ 
10    for  $j \in [0, |R_i| - \mu)$  do // fill up next population with individuals from next
    rank
11       $P^{(k)} \leftarrow P^{(k)} \cup \{R_{i,j}\}$ 
    //  $R_{i,j}$  denotes the  $j$ th individual from rank  $i$ , the rank is assumed to be
    sorted according to the density estimator as explained above; i.e.,
     $j < k \implies d(R_{i,j}) < d(R_{i,k})$ 
12   $i \leftarrow i + 1$ 

```

The original NSGA-II algorithm uses a binary or real-coded search space representation with respective standard crossover and mutation operators. Traditionally, binary tournament selection (see Section 4.5.5) is employed, where the *crowding comparison operator* determines the winner of the competition. The total computational complexity of NSGA-II is bounded by the non-dominated sorting algorithm, i.e., $O(MN^2)$ in the original implementation.

4.11.4 Strength Pareto Evolutionary Algorithm

Like NSGA-II, the SPEA is a multi-objective GA aimed at approximating the Pareto-optimal set of solutions (Zitzler and Thiele, 1998a). Its flow is almost identical to traditional genetic algorithms. That is, a random initial population is generated and evolved through selection, recombination and crossover. The main differences compared to both single-objective GAs and related multi-objective approaches is the determination of the fitness of candidate solutions and the use of an external archive in which the Pareto-optimal individuals—with respect to the explored search space—are stored. The archive constitutes an elitist selection mechanism, in that Pareto-optimal solutions are preserved regardless of evolution effects. In addition it serves as the fitness metric for the GA population, since the merit of

individuals is determined by relating their position in the objective space to the Pareto front approximation in the archive. Finally, because both the population and the archive are considered in the selection process, the algorithm can also be considered a two-population approach, with an indirect seeding-like solution exchange.

In several empirical studies, SPEA could be demonstrated to successfully solve optimization problems with inherent multi-objective properties, i.e., conflicting objectives. Especially an improved version, termed SPEA2 (Zitzler et al., 2001), which incorporates a fitness measure that puts a stronger emphasis on the diversity of (dominated) solutions and a slightly modified routine for the composition of the Pareto set archive, showed good results for a number of benchmarks problems (Zitzler et al., 2001). More importantly, it has been successfully applied to a number of real-world problems.

The fitness assignment procedure of SPEA is schematically demonstrated in Algorithm 4.7 and can be summarized as follows: Prior to the actual fitness assignment, the Pareto archive A is expanded by non-dominated candidate solutions of the population. The archive is then optionally reduced to a predefined maximum size. This reduction routine, for which details will be presented later in this section, is performed in order to prevent a genetic drift that could be caused by highly dense portions of the Pareto front, but also for feasibility reasons, for example, to account for limited storage capacity or post-optimization selection routines. After the archive has been so conditioned, each individual in the archive is tagged with a value, termed strength ($S(\cdot) \in [0, 1)$), reflecting the amount of individuals in the population it currently dominates (l. 5). The strength value constitutes the (reciprocal) fitness of the Pareto-optimal solutions. In contrast, the (reciprocal) fitness of population solutions is given as the sum of the strength values of all Pareto solutions by which they are dominated, increased by a penalty of one in order for Pareto-optimal solutions to receive a higher selection probability (l. 7).

Algorithm 4.7: SPEA Fitness Assignment

```

Data:  $A, P$  // archive and population
Data:  $s_{\max}$  // maximum archive size
1  $A \leftarrow A \cup \{\mathbf{p} \in P \mid (\nexists \mathbf{a} \in A \cup P)(\mathbf{a} \succ \mathbf{p})\}$ 
2 while  $|A| > s_{\max}$  do
3    $A \leftarrow \text{reduce}(A)$ 
4 foreach  $\mathbf{a} \in A$  do
5   
$$S(\mathbf{a}) \leftarrow \frac{|\{\mathbf{p} \in P \mid \mathbf{a} \succ \mathbf{p}\}|}{|P| + 1} \quad // \text{ strength}$$

6 foreach  $\mathbf{p} \in P$  do
7   
$$R(\mathbf{p}) \leftarrow 1 + \sum_{\mathbf{a} \in A_{\mathbf{p}}} S(\mathbf{a}) \quad A_{\mathbf{p}} := \{\mathbf{a} \in A \mid \mathbf{a} \succ \mathbf{p}\} \quad // \text{ (raw) fitness}$$

```

The example in Figure 4.17 illustrates the result of the fitness assignment routine. The Pareto-optimal individuals (marked with crosses) are ranked by their strength values, that is, by the number of individuals they dominate. Larger values mean lower fitness. Thus archive solutions dominating a large number of individuals will be selected with a slightly smaller proportion than solutions that “cover” only few candidate solutions. Broadly speaking, this scheme promotes the exploration behavior of the algorithm, as it leads to a preference for sparsely visited regions of the objective space. A similar property is also given by the fitness of the remaining individuals: They are discriminated not only by the number of dominating archive solutions, but also by the collective strength that these Pareto-optimal candidates impose on them. While the former criterion can be considered as a classical ranking

mechanism, the latter exhibits similarities with sharing, which, as shown above, is a niching technique used in NSGA. In order to avoid an over-representation of similar or equal traits, individuals situated in a denser niche receive a fitness penalty. In Figure 4.17 niches are indicated by differently shaded rectangles. To all individuals in one niche the same fitness is assigned. Later, Zitzler et al. (2001)




Figure 4.17: Strength and fitness assignment in SPEA: In this example (after Zitzler and Thiele, 1998a), the Pareto-optimal set for a given iteration step consists of four candidate solutions (marked with crosses \times), each of which is assigned a strength value to. The strength is the ratio of the number of solutions dominated by the Pareto-optimal individual to be evaluated and the total number of candidates in the dominated population (increased by one). The strength value constitutes the (reciprocal) fitness of Pareto set individuals. The (reciprocal) fitness of population solutions (o) is given as the sum of the strength values of all Pareto solutions by which they are dominated, increased by a penalty of one in order for Pareto-optimal solutions to receive a higher selection probability.

considered this strategy problematic, as they empirically observed a degenerated convergence behavior within the niches, since there solutions exhibit an identical fitness value. The fitness assignment routine was hence been altered in the following way: The strength assignment is extended to also incorporate individuals of the GA population. Thus,

$$S(\mathbf{p}) := |\{\mathbf{q} \in P \cup A \mid \mathbf{p} \succ \mathbf{q}\}|.$$

is evaluated for all individuals both in the GA population and the archive. Note that the normalization to the population size is omitted compared to the original SPEA implementation. For the modified raw fitness we have that

$$R(\mathbf{p}) := \sum_{\mathbf{p}' \in P'} S(\mathbf{p}') \quad P' := |\{\mathbf{q} \in P \cup A \mid \mathbf{q} \succ \mathbf{p}\}|.$$

This scheme provides an improved means to balance exploration and exploitation: The sharing property from SPEA is maintained, since solutions dominated by “strong” individuals, i.e., individuals that dominate a large number of other solutions, still receive a higher penalty. In contrast to the original implementation, however, solutions in the so-defined niches are distinguishable as their mutual domination relation is taken into account.

In addition to the implicit sharing mechanism explained above, an explicit objective-space diversity metric, based on the k -nearest neighbor (k NN) density estimate (see, for example, Devroye and Wagner, 1977), is also incorporated. The density of the surrounding of an individual \mathbf{p} can be estimated by considering the volume $V_{k,\mathbf{p}}$ of a sphere that is centered at \mathbf{p} and contains k samples. A common density function f of the population can then be approximated in the vicinity of \mathbf{p} by

$$f(\mathbf{p}) \approx f_n(\mathbf{p}) = \frac{k}{nV_{k,\mathbf{p}}},$$

where n is the population size. Zitzler et al. (2001) proposes a simplification of the k NN method, where the density estimate is given by

$$D(\mathbf{p}) = \frac{1}{2 + r_k},$$

where r_k is the radius of a hypersphere that is centered at \mathbf{p} and has a minimum volume that contains k elements (schematically demonstrated in Figure 4.18). The selection of the parameter k is critical, in that it determines the range in which the density is approximated. In following Silverman (1986), Zitzler et al. (2001) propose to use the square root of the number of samples, which in this case is the sum of the archive and the population.




Figure 4.18: The raw fitness of SPEA is supplemented by an explicit density approximation, based on a k nearest neighbor variant. The estimate of the density in the vicinity of a candidate solution \mathbf{p} is given by the radius of a hypersphere that is centered on that solution and encompasses exactly k other individuals.

As mentioned before, the external Pareto-optimal set is limited to a predefined size. This is done for mainly three reasons: (1) to limit the space complexity of the algorithm, (2) to ensure a manageable Pareto front size for post-optimization selection procedures, and (3) to prevent the algorithm to prematurely converge because of a genetic drift due to an over-accentuation of portions of the Pareto front. The authors therefore propose a reduction scheme that is concerned with an equitable sampling of the Pareto-optimal set, called average linkage method.⁵⁰ The routine is presented in Algorithm 4.8 and can be described as follows: If the Pareto-optimal archive contains more individuals than a predefined limit, it is reconstructed by clustering individuals on the Pareto front and by combining the representatives of these clusters to form the new approximation of the Pareto front. Clusters are generated iteratively.

⁵⁰The authors refer to the studies of Morse (1980) as their motivation for this procedure.

The initial cluster set consists of sets containing the original archive members. Clusters whose elements exhibit a minimum average distance to each other are collapsed into one set. This procedure is repeated until the cardinality of the cluster set equals the maximum archive size. Afterward, the clusters are replaced by their representatives, for example, their centroids.

Algorithm 4.8: SPEA Archive Reduction

```

Data:  $A$  // archive
Data:  $s_{\max}$  // maximum archive size
1  $\mathcal{C} \leftarrow \bigcup_{a \in A} \{a\}$  // init cluster set; each cluster consists of one Pareto-optimal
   solution
2 while  $|\mathcal{C}| > s_{\max}$  do
3    $J \leftarrow \{1, \dots, |\mathcal{C}|\}^2 \setminus \{(i', j') \mid i' = j' \vee (j', i') \in \{1, \dots, |\mathcal{C}|\}^2\}$  // index set
4    $(i, j) \leftarrow \arg \min_{(i, j) \in J} (\text{avgDistance}(C_i, C_j))$  // get indices minimizing the average
   distance between two clusters
5    $\mathcal{C} \leftarrow (\mathcal{C} \setminus \{C_i, C_j\}) \cup \{C_i \cup C_j\}$  // remove original clusters and add combined cluster
6  $A \leftarrow \emptyset$ 
   // reconstruct archive such that each element represents one of the previously
   computed cluster
7 foreach  $C \in \mathcal{C}$  do
8    $A \leftarrow A \cup \{p_C \mid p_C = \text{centroid}(C)\}$ 

```

In the second version of SPEA, the authors refined both the Pareto archive and the truncation process. In contrast to its predecessor, SPEA2 uses a fixed archive size. Therefore if, after the archive update step in which the non-dominated solutions are copied into the archive, the archive size is smaller than its predefined maximum value, the best dominated solutions are also copied into the archive such that it is completely filled. Allowing well performing individuals from the set of dominated solutions to survive may be suited to increase the convergence toward the Pareto front in the case of too high a diversity, that is, if the algorithm has yet to find a viable search direction. In the other case, if the cardinality of the archive exceeds the maximum size, an archive reduction is performed, similarly to the routine in the original SPEA implementation. Instead of employing a clustering technique, however, the authors developed a different truncation routine in which Pareto-optimal solutions whose objective vectors exhibit the smallest distance to neighbors are iteratively removed. This procedure is less prone to removing individuals from the extremes of the Pareto front than the original clustering method. But it is also rather expensive in terms of computation time ($O(N^3)$). In a later publication (Zitzler et al., 2004), the authors propose a more efficient, albeit non-deterministic algorithm for this task. Hiroyasu et al. (2005) proposed the introduction of measures to increase the decision space diversity by employing mating restrictions and an additional archive. The authors applied their SPEA2+ approach to a tri-objective engine optimization problem, for which a slightly improved spread of the Pareto front was obtained compared with SPEA2 (and NSGA-II). Another extension aiming at improving the selection performance of candidates from the external archive, using a quad-tree representation, was proposed by Mostaghim and Teich (2005). Test results on a number of benchmark functions, however, showed that the performance increase was limited to situations of small archives and large populations, which the authors associate with optimization problems with high numbers of objectives.

The performance of SPEA2 has been tested with a large number of well-established and newly developed benchmark problems. It has been demonstrated to outperform its predecessor SPEA and to exhibit a slightly more homogeneous sampling of the Pareto front than competing approaches such as NSGA-II (Zitzler et al., 2001).

Table 4.1: Domination relations on individual objective vectors and corresponding Pareto set approximation interpretations (reproduced from Zitzler et al., 2003).

relation	object vectors	Pareto set approximations
strict dominance	$\mathbf{a} \succ \mathbf{b} \quad (\forall i)(a_i > b_i)$	$\mathfrak{A} \succ \mathfrak{B} \quad (\forall \mathbf{b} \in \mathfrak{B})(\exists \mathbf{a} \in \mathfrak{A})(\mathbf{a} \succ \mathbf{b})$
dominance	$\mathbf{a} \succ \mathbf{b} \quad (\forall i)(a_i \geq b_i) \wedge (\exists i)(a_i > b_i)$	$\mathfrak{A} \succ \mathfrak{B} \quad (\forall \mathbf{b} \in \mathfrak{B})(\exists \mathbf{a} \in \mathfrak{A})(\mathbf{a} \succ \mathbf{b})$
weak dominance	$\mathbf{a} \succeq \mathbf{b} \quad (\forall i)(a_i \geq b_i)$	$\mathfrak{A} \succeq \mathfrak{B} \quad (\forall \mathbf{b} \in \mathfrak{B})(\exists \mathbf{a} \in \mathfrak{A})(\mathbf{a} \succeq \mathbf{b})$
outperformance		$\mathfrak{A} \triangleright \mathfrak{B} \quad \mathfrak{A} \succeq \mathfrak{B} \wedge \neg(\mathfrak{A} = \mathfrak{B})$
incommensurability	$\mathbf{a} \parallel \mathbf{b} \quad (\exists i)(a_i \geq b_i) \wedge (\exists i)(b_i \geq a_i)$	$\mathfrak{A} \parallel \mathfrak{B} \quad \neg(\mathfrak{A} \succeq \mathfrak{B}) \wedge \neg(\mathfrak{B} \succeq \mathfrak{A})$

4.11.5 Performance Assessment

Similarly to the single-objective case, the performance of MOEAs is often assessed empirically, using benchmark functions. The analysis of the so-obtained results, however, is more complex, since it has to capture the quality of the Pareto set approximation, such as its proximity to the actual Pareto front and the sampling properties. For that purpose a number of quality indicators have been proposed, which are aimed at assessing either a single Pareto set approximation (unary indicators) or at comparing two (binary) or more (multivariate) approximations, where binary approximations can be generally attributed with a viable compromise between complexity and inferential power (Zitzler et al., 2003).

In general, a quality indicator is defined as a function whose arguments are one or more Pareto set approximations and which returns a scalar quality measure for that input:

$$I : \Omega^n \rightarrow \mathbb{R},$$

where Ω is the set containing all Pareto set approximations. Moreover, in order to compare the performance of two Pareto set approximations, it is convenient to transfer the dominance relations to the domain of Pareto sets (see Table 4.1, taken from Zitzler et al. (2003)). There, an additional relation, termed “better,” is introduced, which is defined as the weak dominance of one approximate by a second non-identical set. The utility of a quality indicator or a set of indicators can hence be defined as its interpretation power in terms of these relations.

Numerous indicators have been proposed, the most common ones include:

- the unary and binary *hypervolume indicators*,
- the *set coverage indicator* and
- the binary (*additive*) *epsilon operator*.

The unary *hypervolume indicator* was originally proposed by Zitzler and Thiele (1998b). It measures the volume defined by the Pareto approximation in which solutions are dominated. For it, a reference point (\mathbf{w}) out of the dominated set has to be determined. For each candidate in the Pareto approximation set (\mathfrak{A}), the hyperrectangle spawned by the solution and the reference point ($\square_{\mathbf{aw}}$) is computed. The union of all hyperrectangles determines the covered hypervolume:

$$I_H(\mathfrak{A}) := \text{volume} \left(\bigcup_{\mathbf{a} \in \mathfrak{A}} \square_{\mathbf{aw}} \right). \quad (4.15)$$

This indicator is said to be \triangleright -complete (Zitzler et al., 2003),⁵¹

$$\mathfrak{A} \triangleright \mathfrak{B} \implies I_H(\mathfrak{A}) > I_H(\mathfrak{B}).$$

Furthermore, it is $\not\prec$ -compatible (ibid.),

$$I_H(\mathfrak{A}) > I_H(\mathfrak{B}) \implies \mathfrak{B} \not\prec \mathfrak{A}$$

⁵¹A similar order-homomorphism, termed *Pareto dominance compliance*, was proposed by Knowles and Corne (2002).

The hypervolume indicator can be extended to a binary metric in a straightforward manner, for example, as follows Zitzler (1999):

$$I_{HD}(\mathfrak{A}, \mathfrak{B}) := \begin{cases} I_H(\mathfrak{A}) - I_H(\mathfrak{B}) & \mathfrak{B} \triangleright \mathfrak{A} \\ I_H(\mathfrak{A} \cup \mathfrak{B}) - I_H(\mathfrak{B}) & \text{otherwise.} \end{cases}$$

This metric yields the hypervolume covered by \mathfrak{A} and not \mathfrak{B} . In case \mathfrak{B} weakly dominates \mathfrak{A} , for Pareto dominance compliance, the volume difference is obtained. As noted by Van Veldhuizen (1999), the hypervolume metric may be problematic for non-convex Pareto fronts. In addition, it strongly depends on the selected reference point.

The *set coverage indicator* (Zitzler and Thiele, 1998b) determines the ratio of elements in \mathfrak{B} that are (weakly) dominated by elements of \mathfrak{A} . Formally, we have that,

$$I_C(\mathfrak{A}, \mathfrak{B}) := \frac{|\{\mathbf{b} \in \mathfrak{B} | (\exists \mathbf{a} \in \mathfrak{A})(\mathbf{a} \succeq \mathbf{b})\}|}{|\mathfrak{B}|}. \quad (4.16)$$

Like the set coverage indicator, the *additive epsilon indicator* (Zitzler et al., 2003) is Pareto dominance compliant. Given two Pareto set approximations $\mathfrak{A}, \mathfrak{B}$, the indicator I is defined as the minimum epsilon value (or translation) for which \mathfrak{A} (weakly) dominates \mathfrak{B} ,

$$I(\mathfrak{A}, \mathfrak{B}) := \arg \min_{\epsilon} \{ \epsilon \in \mathbb{R} | (\forall \mathbf{b} \in \mathfrak{B})(\exists \mathbf{a} \in \mathfrak{A})(\mathbf{a} \succeq_{\epsilon} \mathbf{b}) \}. \quad (4.17)$$

A large number of other indicators have been proposed,⁵² for example, to measure the sampling behavior of the Pareto set approximation (Schott, 1995; Van Veldhuizen, 1999; Zitzler, 1999; Deb et al., 2002a), to estimate the ratio of final points that are not members of the actual Pareto front (Van Veldhuizen, 1999), evaluate the average distance of candidates from the Pareto front (Van Veldhuizen, 1999), or metrics based on utility functions (Hansen and Jaszkiewicz, 1998). Typically, not a single indicator but a combination of different indicators is used (for example, see Zitzler et al., 2003).

A Pareto set approximation resulted from a non-deterministic optimization approach such as a MOEA is random, reflecting a representative of a specific probability distribution (e.g., see Knowles et al., 2006). Thus in order to infer significant statements about the performance of an approach with respect to a specific problem, this distribution has to be estimated. Using, for example, an indicator-based comparison, the desired information is if one algorithm yields a “better” distribution of indicator results for a problem than another one, for which several non-parametric statistical distribution tests have been proposed (ibid.).

Empirical tests can only be aimed at comparing the performance of optimizers on specific problems. This approach does not allow for strict general performance statements about the studied optimizer. Corne and Knowles (2003a) showed that despite the differences in the performance measure (broadly speaking, Pareto optimality instead of scalar fitness value) the NFL, which assumes a single-objective approach, also holds in the multi-objective case. Moreover, the same authors (ibid.) proved that the NFL is applicable also for sub-problems that can be associated with particular shapes of the Pareto front, such as high level of discontinuities. The authors argue, however, that there are free lunches, which they term “free leftovers,” when comparing two algorithms with regards to specific quality measures. Another potential for free lunches can be attributed to elitism (Corne and Knowles, 2003b): Since different approaches employ different archiving techniques in order to collect intermediately Pareto-optimal individuals, general performance differences between alternative algorithms may occur.

Because of the high attention to multi-objective optimization approaches, especially in the past decade, but also because of the different foci when comparing their performance, a multitude of benchmark functions and test suites have been proposed. Comprehensive discussions and surveys were, for example, presented by Van Veldhuizen and Lamont (1998, pp. 61–71), Deb (2001, pp. 338–389) and Coello Coello et al. (2007, chapt. 4). As a baseline for the conception of multi-criteria test functions,

⁵²Comprehensive discussions were presented by Van Veldhuizen (1999) Zitzler (1999) Deb (2001, pp. 324–333) and Coello Coello et al. (2007, chapt. 5).

similar aspects as in the single-objective are to be considered (Deb, 1999). For example, multi-modality, impaired building block processing (deception or epistasis), isolated optima and noise play an important role for the convergence toward the final Pareto set approximation. Additionally, however, other factors, including a non-convex or discontinuous Pareto front, other discontinuities in the search space, or non-uniform candidate distributions may have a strong impact on both the convergence and the properties of the final Pareto front approximation. On that score, Deb (1999) presented a number of guidelines for the construction of benchmark problems in order to capture important issues and goals for multi-objective optimization approaches. Similar considerations have led to other comprehensive test suites (for example, Zitzler et al., 2000; Deb et al., 2002b, 2006).

Another common approach used in the empirical assessment of the performance of probabilistic algorithms on benchmark functions is the attainment indicator. It states whether a specific goal was achieved in a single run of a specific non-deterministic multi-objective optimizer (Fonseca and Fleming, 1996; Grunert da Fonseca et al., 2001). Given a Pareto set approximation \mathfrak{A} out of all possible approximations Ω , the attainment indicator is given by the probability that one of the solutions (weakly) dominates the goal \mathbf{z} :

$$\alpha(\mathfrak{A}, \mathbf{z}) := P((\exists \mathbf{a} \in \mathfrak{A})(\mathbf{a} \succeq \mathbf{z}))$$

The empirical estimate over a number of runs (n_r) can hence be stated as follows:

$$\alpha(n_r, \mathbf{z}) := \frac{1}{n_r} \sum_{i=1}^{n_r} \mathbf{1}_{\succeq}(\mathfrak{A}_i, \mathbf{z})$$

where \mathfrak{A}_i is the Pareto set approximation after run i and \succeq denotes the (weak) domination set relation (i.e., at least one of the solutions dominates \mathbf{z}), of which $\mathbf{1}_{\succeq}$ is the indicator function. In addition to an estimate of attainable goals, this formulation allows for a number of distribution-related analysis, for example, to determine fundamental sampling differences between two algorithms. As pointed out by Grunert da Fonseca et al. (2001), the attainment function corresponds to a first-order moment measure. Thus in order to infer and compare more critical aspects of attainable Pareto approximation distributions, they propose to extend the analysis to higher-order moments.

In addition to forming the basis for the quality assessment of different multi-objective optimization approaches, recent research activities propose the direct incorporation of such metrics into the selection schemes of the algorithms itself. Especially the epsilon and the hypervolume indicators have been reported to yield highly competitive results (Zitzler and Künzli, 2004; Beume et al., 2007; Igel et al., 2007; Brockhoff et al., 2008; Bader and Zitzler, 2011). Traditional approaches implicitly bias the convergence toward the Pareto set approximation by the rank-based selection and the diversity-preserving measures. The rationale behind incorporating quality measures into the search direction is to replace this potentially unintentional bias by a quality-measure-driven one. Alternatively, such a method also allows for the inclusion of metrics specified by the decision maker. One of the main drawbacks of these recent developments is in turn their risk of yielding approximations that are exceedingly well adapted to the incorporated quality measure but exhibit a less beneficial performance in terms of other metrics. Another means to circumvent algorithm-induced bias is the explicit inclusion of DM preference (e.g., see Fonseca and Fleming, 1993; Greenwood and Hu, 1997; Branke et al., 2001; Phelps and Köksalan, 2003; Deb and Sundar, 2006; Jaszkievicz, 2007; Thiele et al., 2009; Greco et al., 2010).

As argued before, the selection of an optimal MOEA is impossible, since the performance of each of the approaches is strongly problem-dependent. Algorithms “outperforming” others on a large number of test problems may be considered preferable. Again, however, the performance measures and the test methodology strongly bias the outcome of any comparison. Along that line, we have based our choice to use two state-of-the-art MOEAs—NSGA-II and SPEA 2—on a number of competitive comparison results (e.g., Zitzler and Thiele, 1998b; Zitzler et al., 2000, 2001; Deb et al., 2002a; Khare et al., 2003), and not on the assumption of using an optimal approach (Seifert, 2006b). The complexity and the high computation times of the problems aimed at in this work can be considered prohibitive to serve as a platform to extensively compare numerous approaches (ibid.). We have not included quality-measure-driven or preference-incorporating approaches into our consideration.

5 Hybrid Local/Meta-heuristic Optimization Approach

As we have shown in the previous chapter, evolutionary algorithms are effective global optimization approaches that, when tailored accordingly, are suited for a wide range of applications. We have also discussed their limitations: In particular, the ones related to the search space representation, constraint handling or multi-modality of the search space. And more generally, by following the argument of the no free lunch (NFL) and related theories, in acknowledging that evolutionary algorithms—like any other optimization algorithm—cannot be considered universally ideal. An ideal approach would indeed involve the selection of a search strategy that best fits the search space, which in many applications little a priori knowledge is available of. Combining different optimization routines to be prepared for various problem characteristics thus appears to suggest itself. A common and naive but yet not necessarily ineffective approach is to subsequently carry out different search algorithms. For example, the final result of a genetic algorithm may serve as an initial value of a consecutively performed local optimization.¹ Along that line, more involved methods, with a tighter coupling of different biology-inspired and local search (LS) operations have been proposed (for example, Brady, 1985; Ackley, 1987a; Suh and Van Gucht, 1987; Mühlenbein, 1989; Jog et al., 1989).² Those approaches using an evolutionary algorithm are commonly referred to as genetic local search.

5.1 Related Work on Memetic Algorithms

The term *memetic algorithm* was coined by Moscato (1989), adopting the concept of “*memes*,” which, broadly put, are defined as transmittable quasi-entities carrying cultural properties such as capabilities or conceptions. They have been first introduced as an analogy to genes by Dawkins (1976).³ Moscato (1989) was among the first to acknowledge that genetic local search approaches are to be regarded as a new class of algorithms rather than just the combination of two or more optimizers. Consequently, he stressed the importance to account for the interaction of the involved global and local search components, for example, by tailoring genetic operators.

The socio-biological perspective on EA search hybrids due to his and previous work by others (most prominently, by Hinton, 1987) led to further discussions, aiming to clarify which natural role-model the hybrid approaches should follow (see, for example, Turney, 1996). Specifically, the question whether algorithms should rather follow a *Baldwinian* or a *Lamarckian* learning scheme was frequently raised.

In *Baldwinian evolution* (Baldwin, 1896), learning can be regarded as a second-order mechanism: Not the phenotype itself but its interaction with the environment determines the selection pressure on individuals (Smith, 1987; Dennett, 2003). It is often associated with a degree of learning or adaptation capacity of individuals. In the scope of memetic algorithms, this can be broadly attributed to a trait that can be described as how well-suited an individual is in terms of serving as a start solution for a

¹Another example for search compositions is the iterative or concurrent start of multiple instances of local optimizers (see, for example, Lourenço et al., 2010; Martí et al., 2010), which is not discussed here.

²Other examples are discussed in Davis (1991).

³In his book “The Selfish Gene,” Dawkins (1976) precautionarily disclaims that he imputes a will to genes. Still, he transmogrifies genetic material into the driving force, the subject not only of evolution but also of individuals. Consequently, turning everything upside down, Dawkins translates his concepts of an invisible force to the domain of individual-environment interaction (ibid.). Despite our opposition to this theory, we consider the concept of “memes,” entities carrying programs with the potential to improve candidate solutions, as quite adequate for the description of genetic local search approaches. We will hence borrow the term *memetic algorithm* from Moscato (1989), albeit ignoring his martial art analogy.

local refinement. It is clear that this property strongly depends on the employed local optimizer, its control parameters and on the representation mapping between genotype, phenotype and expressed phenotype for the local optimizer. The Baldwin effect does not found a principally different theory of evolution, but rather constitutes an elaboration on the cause of selection pressure. It can hence be said to be “compatible” with Darwinian evolution. This is also reflected when designing Baldwinian memetic algorithms, which are sometimes also referred to as cultural algorithms. After obtaining the phenotype of a candidate solution, it is mapped into a local search representation, local search is conducted, and the obtained fitness also represents the fitness of the original individual. No genotypic modification occurs. This property imposes a number of difficulties on the MA model: (1) Using non-deterministic local optimizers also renders the fitness evaluation non-deterministic: The same genotype may lead to significantly different pay-offs, and different genotypes may lead to the same pay-off. (2) Local search information obtained through the local improvement step cannot be directly exploited in the evolutionary search and is hence not reflected in the information processing of genetic algorithms, impairing its mixing behavior.

In contrast, *Lamarckian evolution* assumes a direct interaction of the environment with the genotype. This property is often illustrated with Lamarck’s false “explanation” of the giraffes’ elongated necks and forelegs, which he attributed to a materialization of their constant habit to reach for high-hanging leaves. Albeit Lamarckian evolution is claimed to be the underlying role-model for many memetic algorithms, this analogy comprises a number of ambiguities. Turney (1996), for example, argues that strictly with a Lamarckian approach, local search would have to be performed on the same (genotypic) domain as the evolutionary algorithm; or at least the expression operator transcoding the genotype into the phenotype would have to be invertible in order to directly incorporate the local improvement.⁴ Moreover, while the incorporation of local-search-obtained properties into the genotype may be favorable for the EA in that it is then implicated into the course of exploration, it is at the same time the source of potentially strong schema disruptions (ibid.). Finally, too high a local search intensity, for example, by triggering the local search too often, may degrade the diversity of the population (Orvosh and Davis, 1994).

Besides a categorization according to the biological analogy, several attempts have been made to formalize memetic algorithms. Radcliffe and Surry (1994), for example, promoted a uniform framework for genetic and memetic algorithms, albeit restricted to less involved interactions of the constituting search components. Their concept, however, modeling local search as a unary modification operator, similar to, say, mutation, renders it a concise and straightforward extension of EAs.

Since the work of the previously mentioned authors—especially the technical report of Moscato (1989) gained significant attention—a large number of efforts in the field of memetic computation have been exercised (for an overview, see Hart et al., 2005).⁵ In the following, we will summarize a number of related studies that we have based our considerations upon.⁶

Freisleben and Merz (1996) applied a hybrid of a local search heuristic and a genetic algorithm to combinatorial (traveling salesman) problems. In their approach, all individuals selected for recombination are also subjected to the local search. The reported results are considerably superior compared with those obtained through a plain GA. The authors stress the importance of accounting for problem-specific properties, for example, by tailoring crossover operators. Later (Merz, 1997), the method was used on a quadratic assignment problems, performing similarly well.

Hart (1994) proposed a memetic algorithm in which the local search frequency is adaptively determined, for example, by accounting for the distance of solutions to previously discovered local optima. For the numerical examples used by Hart (Griewank’s and Rastrigin’s functions, see Appendix F), applying local search with low frequencies showed the best results both in terms of convergence rate and success in finding a solution close to the exact local optimum. Land (1998) extended this concept to combinatorial problems, where he found that Baldwinian is as effective as Lamarckian learning, a conclusion that is in contrast to Hart’s findings. Land also stresses the importance to consider the search intensity. He distinguishes between two extremes: (illustrated in Figure 5.1). (1) greedy search

⁴We shall see that a similar situation also has to be solved for in this work.

⁵Another noteworthy direction is the attempt to hybridize multi-objective evolutionary algorithms (Goh et al., 2009).

⁶A comprehensive survey on related (adaptive) memetic algorithms was given by Ong et al. (2006).

of the local environment until no further improvement is found or (2) a slack search that stops as soon as an improvement of the original candidates is found. Land argues that an adequate choice is strongly problem-dependent and should hence be included in adaptive MAs, a theme that was later seized and developed by other authors.




Figure 5.1: Local search intensity in memetic algorithms: Given a maximization problem and a candidate solution A obtained through an evolutionary step; assuming that a local search approach would yield iterates B – D , a greedy local search would not stop until the final local optimum D is achieved, whereas a slack variant would yield B since it is the next solution that is better than A .

Krasnogor (2002) developed an adaptive MA, termed multi-meme algorithm that uses a set of competing local search approaches. For each candidate solution to be locally improved, the appropriate local optimizer is determined by means of inheritance from the parent solutions. The inherited meme is stored alongside with its application probability. This scheme is broadly inspired by the meme transmission model of Dawkins (1976). Competitive results for a number of combinatorial problems were demonstrated. In addition to this adaptive meme selection and transmission scheme, Krasnogor (2002) proposed a concept of a self-generating MA, in which memes are co-evolved, either using parameters influencing the properties of the local searchers or using a genetic programming approach to rigorously evolve local search algorithms. In a proof-of-concept example, Krasnogor and Gustafson (2004) successfully apply their approach to hard combinatorial problems. There, a specific emphasis is put on the “move” operator, which determines the neighborhood for the local search and whose co-evolution is conjectured to be of major importance since it should be adapted to the (generally unknown) fitness landscape. The authors claim that their concept could also be used for other tasks. Further studies are needed, however, to assess how a co-evolved meme would compare to sophisticated numerical search routines.

Another self-adaptive system was proposed by Smith (2007), termed coevolving memetic algorithm (COMA), that makes use of a separate population containing representations of the local search approaches supplemented with configuration parameters. The payoff of a meme is given by the fitness improvement that it has achieved. The selection scheme of memes is itself a parameter subject to evolution: It can be either conducted in a fitness proportionate fashion, similarly to Ong and Keane (2004), fully randomly or tightly linked to the solutions. Being aimed at combinatorial problems, only simple local optimizers (greedy and steepest ascend) were employed. Their search intensity configuration is subject to coevolution. The feasibility of the approach has been demonstrated on several combinatorial test problems, albeit yielding ambiguous performance results compared with a non-adaptive MA and a GA.

Cowling et al. (2001) developed an adaptive MA that they termed hyperheuristic. The basic idea is to provide the algorithm with a number of low-level heuristics from which for each the appropriate routine is selected by the hyperheuristic. Different selection mechanisms were studied, including a random selection and selection schemes based on the observed improvement capabilities of the memes (low-level heuristics). Competitive results were achieved for a number of (scheduling) problems (also see Kendall et al., 2002). The authors emphasize the ability of their approach to adaptively generate adequate heuristics through composition of relatively simplistic memes.

Ong and Keane (2004) propose what they call a meta-Lamarckian learning procedure. Prior to a local improvement step, a search method is adaptively selected out of a user-provided pool of numerical search algorithms. Two selection schemes were proposed:

- (1) *Sub-problem decomposition*. Each locally modified candidate is stored along with the LS that was applied and a reward measuring the improvement due to the LS. In subsequent steps, the best performing LS (according to the reward) out of the k nearest neighbors of the candidate to be improved is selected. The rationale behind this approach is that search spaces often exhibit differently conditioned regions on which different LS perform best.
- (2) *Biased roulette wheel selection*. A global approach in which the local optimizer is chosen through roulette wheel selection, where the wheel segments reflect the reward measure of the local search approaches.

For numerical experiments and one structural optimization problem, their approach yielded highly competitive results compared with a canonical GA and a multitude of different MAs using different local search approaches. Interestingly, the global, roulette wheel local search selection clearly outperformed the sub-problem decomposition one.

5.2 Domain of Local Search Operations

In this work, we employ a perspective similar to the formal memetic algorithm description by Radcliffe and Surry (1994) in that we regard the local improvement step as a mutation-like operator, rendering our approach a Lamarckian learning variant. In order to discuss the impact of this view, we shall refine some of the assumptions that we have previously omitted for the sake of simplicity. Especially, we have not rigorously distinguished between the terms *fitness*, *merit* and *objective* function. We shall catch up on that in the following. We will restrict our discussion to scalar optimization formulations and, analogously, scalar constraint handling methods. Formally, let the optimization problem be defined as:

$$\begin{aligned}
 & \arg \min_{x_1, x_2, \dots, x_n} f(x_1, x_2, \dots, x_n) \\
 & \text{subject to} \quad x_i^l \leq x_i \leq x_i^u \quad (i = 1, \dots, n) \\
 & \quad \quad \quad \text{other linear or non-linear constraints of the form} \\
 & \quad \quad \quad (g_j(x_1, x_2, \dots, x_n), z_j) \in G_j \subset Z_j \times Z_j \quad (j = 1, \dots, m) \\
 & \text{where} \quad x_i, x_i^l, x_i^u \in X_i \\
 & \quad \quad \quad f : X_1 \times X_2 \times \dots \times X_n \rightarrow \mathbb{R} \\
 & \quad \quad \quad g_j : X_1 \times X_2 \times \dots \times X_n \rightarrow Z_j \\
 & \quad \quad \quad X_i, Z_j \text{ are ordered rings, often } \mathbb{R} \text{ or } \mathbb{Z} \quad (i = 1, \dots, n; j = 1, \dots, m).
 \end{aligned} \tag{5.1}$$

Then, f is called the *objective function*. A *merit function* can be defined as a function that for each parameter input yields the pay-off of the point. The pay-off may be the direct result of the objective function but may also account for constraint violations or apply optimization approach-specific transformations such as a fitness scaling. The merit function ($\hat{\mu}$) thus has the same prototype as f :

$$\hat{\mu} : X_1 \times X_2 \times \dots \times X_n \rightarrow \mathbb{R}.$$

As discussed earlier, GAs typically strictly distinguish between the genotype and the phenotype. The phenotypic representation is then obtained through an expression operator, for example, defined as a function E , with

$$E : X_g \rightarrow X_1 \times X_2 \times \cdots \times X_n,$$

where X_g is the genotype space, in our case $X_g = \{0, 1\}^\ell$, where ℓ is the bit-string length. Consequently, the *fitness function* μ is a composition of expression and merit function:

$$\begin{aligned} \mu : X_g &\rightarrow \mathbb{R} \\ \mu &:= \hat{\mu} \circ E. \end{aligned} \tag{5.2}$$

Genetic operations are performed exclusively on the genotype. Thus if we regard the local improvement as an operator, similar to, say, mutation, we can define it as operator, L with

$$L : X_g \rightarrow X_g.$$

In contrast to mutation, however, which in the binary case consists of only bit operations, the local improvement operator utilizes some model of the merit function. In general, this model is not identical, often not even completely equivalent, to the fitness function of the GA. One reason for the separation is the optimizer-specific domain of the merit function. Local optimization approaches such as the ones used in this work (see Section 5.4), typically assume a continuous search space, and often a smooth function. Moreover, state-of-the-art numerical optimizers further employ tailored, sophisticated constraint handling techniques (cf. Section 5.4.1), which are normally different from the ones employed in GAs. It is thus necessary to maintain multiple merit models and a mapping between the representation of candidate solutions. Let X_l be the domain of the local optimizer. The first step for the local improvement operation is the conversion from the genotypic representation of a candidate solution into the local optimizer's domain:

$$E_l : X_g \rightarrow X_l.$$

Vice versa, the result of the local optimization has to be transcoded back into the GA genotype:

$$E_g : X_l \rightarrow X_g.$$

Often, X_l is equal to $X_1 \times X_2 \times \cdots \times X_n$. Then, the problem definition of (5.1) including the fitness function and constraints, if applicable, can be directly used by the local optimizer. Otherwise a surrogate objective or merit function on X_l has to be defined, analogously to (5.2). In that case, the local optimization is performed on the surrogate function. In Lamarckian learning, as used here, in order for the GA to “take possession” of the local optimization result, the candidate has to be transcoded back to the genotypic representation. In Baldwinian learning, the candidate would stay unaltered and would only take on the fitness value obtained through local search. The merit function domains of the GA and the local optimizer and the conversion operators are summarized in Figure 5.2.

In general, the transcoding operator E_l is not invertible. Moreover, especially in the case of a discrete genotype, like in this work, the result of the merit function on the (continuous) local search domain and the back-coded correspondent may yield different values. Two main problems, which are depicted in Figure 5.4, may arise from this situation:

- (1) Depending on the discretization, small local changes with respect to the domain of the local optimizer may have no impact on the discrete parameter domain of the GA (Dürr, 2006b; Dürr et al., 2007b). In other words, the granularity of the discrete representation may not be able to capture the deviation produced by the local optimization (LO).
- (2) Since the result of the local optimizer can be expected to be more accurate compared to the discrete results obtained with the GA, a general question arises on the distance between the original GA-obtained candidate, the locally optimized version and the corresponding GA representative (Dürr (2006b)).




Figure 5.2: Merit function domains of genetic algorithm and local optimizer and corresponding conversion operators. GA operations are conducted on the genotypic representation (X_g). For the fitness evaluation, the genotype is expressed (E) as phenotype. The local search is in general performed on a different domain (X_l) to which the genotype of the candidate solution that is subject to local modifications has to be transcoded (E_l). If that domain is not identical to the domain of the problem definition, a surrogate merit function, employing a further expression operator (E') has to be introduced. In order to make the locally optimizing result amenable to evolutionary exploration, it has to be coded back to the genotypic domain (E_g).




Figure 5.4: Issues of transcoding between discrete and continuous merit function domains. The dashed line denotes the discrete search space of the GA, the solid line is the graph of the continuous function of the local optimization routine: (a) In a well-conditioned case, the result obtained through local optimization indeed constitutes an improvement of the candidate, also according the genotypic space; (b) in other situations, however, local search may fail to carry the candidate outside of the discretization unit given by the discrete representation: The original and the transcoded locally optimized solution are identical. (c) A more detrimental coding may even lead to a decreased fitness of the back-coded individual.

In addition to the coding frictions that arise from Lamarckian learning approaches,⁷ the local search may drastically impact the schema processing and mixing behavior of the EA. Particularly, similar to other highly disruptive operators (e.g., long-range or high-probability mutation), the EA must be granted enough time to incorporate the local-search generated changes into its evolutionary process. It is thus not surprising that many authors (e.g., Hart, 1994; Orvosh and Davis, 1994) reported on a better performance when applying local search with low frequencies. More rigorously, Radcliffe and Surry (1994), conjecture that memetic algorithm are most successful in cases where the fitness function is decomposable; that is, when small changes of parameters also lead to only small changes of the fitness. This can be translated to fitness landscapes with a low epistasis variance.

Another problem for a memetic algorithm approach is multimodality. For example, the behavior of gradient-based local optimization routines strongly depends on their start value. Small deviation of start values may result a dramatically different local optimum. This local optimum may be in much lesser proximity to the GA-produced start solution than expected or intended. Especially with niching techniques such as restricted tournament selection as used in this work, the local search may be counterproductive in that it drives the candidate solution (which is expected to belong to a certain similarity slot in terms of Hamming distance) away from the designated niche.

5.3 Memetic Algorithm Design

Summarizing the considerations following from the related work discussed above, particularly from the questions raised in the seminal thesis by Hart (1994, pp. 3–5), the most important characteristics of an MA are given by

- the frequency of local search steps,
- the candidate selection for local search,
- the local search intensity and
- the local search efficiency.

In this work, an attempt was made to directly, yet flexibly address these topics. As a result, a scheduler-based system, similar to the rule-based approach proposed by Krasnogor (2002)—which to some extent was realized by Krasnogor and Smith (2000)—was developed (Dürr, 2006b; Dürr et al., 2007a).⁸ This framework assumes the evolutionary algorithm to constitute the core of the optimization procedure. It is supplemented with a number of schedulers that determine when and how to call the local search steps. The schedulers can be categorized as follows:

- *triggers*, determining the frequency of local search calls,
- *candidate selection*, governing the choice of candidates to be locally modified and
- *local research restrictions*, specifying the neighborhood of local search operations.

The framework also provides an entry point for mechanisms to adaptively select the local search approach, albeit in this work only static memes are considered.

The flow is schematically depicted in Figure 5.5. After each GA generation, the trigger function (τ) is called, determining whether a local search step is to be performed or not. In case of a positive trigger result, a subpopulation for the local improvement step is selected. Each candidate of this subpopulation undergoes a local optimization as follows: Optionally, a local search procedure is adaptively selected in the first step. The candidate is transcoded to the local search approach domain (see Section 5.2). Local optimization is performed. The search result is then subjected to model control: If under recoding to the EA domain (genotype), the improved candidate outperforms the original solution, the original individual is replaced by the locally improved one. Otherwise, the LO solution is rejected and the original candidate is kept.

⁷Baldwinian learning does not require any such recoding strategies. Yet, it does not eliminate the problems associated with the different domains of the EA and the local search approach.

⁸Some of its technical details are discussed in Section 6.6.

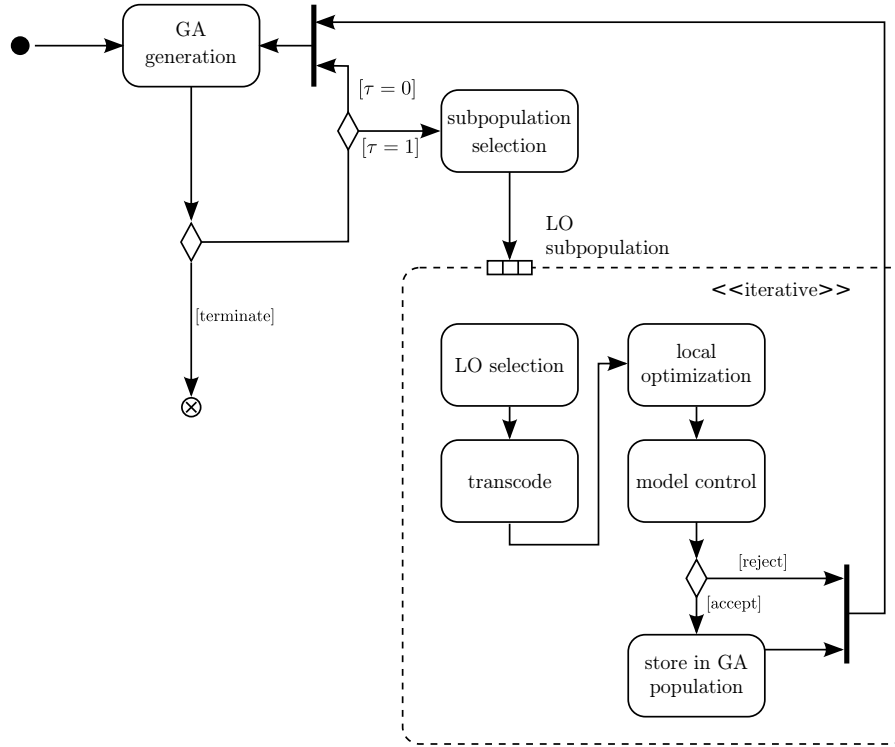


Figure 5.5: Schematic of flow of memetic algorithm.

5.3.1 Trigger

The local search operation, especially when applied as a Lamarckian evolution approach, is closely related to the intensification phase in tabu search (Section 4.1). Like there, intensification has to be applied in a regulated manner, since it is only one side to the story, and diversification is the other. In the evolutionary algorithm jargon, it is the balance between exploitation and exploration. The local search can be expected to put an emphasis on exploitation rather than exploration. In other words, exploitation is accentuated through local search. To allow the GA to reincorporate the locally improved solutions into its exploration process, that is, to allow mixing-in of the obtained local refinements, it has to be granted time. Another argument for a frequency-control of the local search is the computational expenses, which in real-world applications, like the ones shown here, can be expected to be high, in turn also rendering the local search expensive, especially when using first and second-order exploiting routines. We have therefore adopted a flow in which the GA receives the paramount number of function evaluations, and in which the local search is applied after a number of generations. In their memetic algorithm, Areibi and Yang (2004) follow this idea in a straightforward manner, by alternately employing a global and a local search. There, the local search is restricted to a number of iterations, and to a pre-defined radius.

In order to obtain a finer-grained control on when to execute the local search, we have introduced the concept of a trigger function (Dürr, 2006b; Dürr et al., 2007b). The trigger intervals are expressed in terms of GA generations. A trigger function is defined as a decision when to apply local search. It is hence defined as function mapping the current generation index to a decision value (0 or 1):

$$\tau : \mathbb{N} \rightarrow \{0, 1\}$$

In general, the design of this function is strongly problem-specific and can hence be adapted accordingly. In addition, dynamic and self-adapting variants, inferring the convergence characteristics can be realized. Accordingly, three classes of trigger functions can be distinguished.

Static trigger functions constitute the simplest form. A *linear* trigger (Dürr, 2006b; Dürr et al., 2007b), for example, leads to an LO execution every n_g generations:

$$\tau : k \mapsto \begin{cases} 1 & (k+1) \equiv 0 \pmod{n_g} \\ 0 & \text{otherwise,} \end{cases} \quad (5.3)$$

where k is the generation index. This variant is aimed at modeling an alternate execution of the global and local search, similarly to Areibi and Yang (2004). In order to allow for an intensified exploration of the search in early stages, it can be supplemented with a start generation prior to which no local search is allowed.

In order to account for the convergence progress of the hybrid search, *dynamic* triggers can be alternatively applied (Dürr, 2006b; Dürr et al., 2007b). It could be argued, for example, that later search stages should be dedicated to an emphasis on exploitation, assuming that schema mixing and construction do not play a major role in this phase. Under this assumption, an accentuated local search in a later stage could be considered as supporting and intensifying the exploitation behavior of the genetic algorithm, similarly to a cooling scheme in simulated annealing.

As pointed out by Dürr (2006b), however, depending on the problem, opposite arguments may become valid. Similarly to mutation, the local search operator may have a strong disruptive effect on schemata found by the GA, such that its late application impairs the convergence efforts of the GA. Moreover, assuming that the local search operator was successful in obtaining local optima, a repetitive application of the local search on these candidates may not lead to a significant improvement. In such situations, a trigger function that displays its peak in the mid-range in terms of generations may prove better suited.

Analogously to the static trigger function, a dynamic one can hence be defined straightforwardly as:

$$\tau : k \mapsto \begin{cases} 1 & (k+1) \equiv 0 \pmod{T(k)} \\ 0 & \text{otherwise,} \end{cases} \quad (5.4)$$

where $T(k)$ denotes the interval between executions in dependence of the generation index. Alternatively, the dynamic trigger can be defined as a probabilistic function,

$$\tau : k \mapsto \begin{cases} 1 & \mathcal{U} < p(k) \\ 0 & \text{otherwise} \end{cases}$$

where $\mathcal{U} \in [0, 1]$ is a uniform random number and $p(k) \in [0, 1]$ denotes the generation index-dependent trigger probability, for which the following functions have been considered in this work:

- *linear*,

$$p' : k \mapsto p_l + (p_u - p_l)(k - s)\alpha,$$

- *asymptotic*,

$$p' : k \mapsto p_u - \frac{p_u - p_l}{(k - s + 1)^\alpha}, \text{ or}$$

- *exponential*,

$$p' : k \mapsto \begin{cases} p_u & k - s > \alpha \\ p_l + (p_u - p_l) \exp(k - s - \alpha) & \text{otherwise,} \end{cases}$$

where p_l and p_u are the initial and the final probabilities, respectively, s is the minimum generation until which no local improvement operation is triggered, and $\alpha \in [0, 1]$ is user-provided slope parameter. The final probability is then defined as the minimum as the result of the progression function and the maximum probability:

$$p(k) := \min\{p'(k), p_u\}.$$

Adaptive triggers, termed *on-demand* by Dürr (2006b), account for the fitness improvement over a number of past generations. Like the other types of triggers, this variant can be applied in two ways: (1) either assuming that a steady fitness improvement indicates that the previously triggered local search

was successful or (2) that a fitness stagnation prompts for a local search step to prevent genetic drift. Depending on the interpretation, the frequency produced by the trigger is increased or decreased with a certain level of fitness improvement. The fitness metric to be used strongly depends on the optimization problem. While for unimodal functions the fitness improvement may be sufficiently well captured by evaluating only the champion fitness, multimodal functions may require a broader picture, for example, using the average fitness of a portion of the population. Another critical parameter is the number of past generations that are taken into account. Small look-back windows reduce the exploration time given to GA, whereas too large a window may render the trigger virtually inactive. Last, in order to further increase the exploration time for the GA, a relaxation parameter, enforcing a minimum interval between triggers, can be introduced.

$$\tau : k \mapsto \begin{cases} 1 & (\forall k > k' > k - s)(\tau(k') = 0) \wedge (\forall k > k' > k - w + 1)(\Delta \blacktriangleright \Delta\hat{\mu}(k', k' - 1)) \\ 0 & \text{otherwise,} \end{cases} \quad (5.5)$$

where s is the minimum interval between triggers, w is the look-back window—the number of past generations evaluated in terms of fitness improvement— Δ is the required (minimal or maximal) fitness improvement, $\Delta\hat{\mu}(\cdot)$ is the actual fitness improvement between two generations, and \blacktriangleright denotes any valid order relation ($>$, throughout this work).

Finally, the trigger function could take a *self-adaptive* form, for example, by accounting for the success rate of preceding local optimization operations. This scheme, however, has not been investigated here.

5.3.2 Selection

In addition to triggers, determining whether local search is applied at all, the candidates eligible for local improvement have to be drawn. In contrast to GA selection, which aims at an optimum ratio of selection intensity versus productive diversity, the criteria for the local search selection are less obvious. Rationales can be given for virtually any choice (ibid.): below-average performing individuals can be expected to greatly profit from local improvements. Thus selection in a steady-state GA fashion would be advantageous. On the other hand, a selection of well-performing candidates could drastically speed-up the overall convergence, yet potentially causing super-individuals. Also, individuals with an average pay-off can be expected to gain benefit from modifications, as they do exhibit fitness contribution but lack some form of advantage. In all cases, the local search also constitutes the risk of diversity loss. Thus, for example, selecting all individuals out of one category can hence be expected to be “dangerous.”

The first option when selecting individuals to be locally optimized is a *elitarian* selection scheme (Dürr, 2006b; Dürr et al., 2007b), which can take the form of a deterministic or a probabilistic truncation selection. In the former case, a user-provided number (n) of individuals are collected in sequence of their fitness. The dynamic variant is realized in the same fashion as the usual truncation selection operator (Section 4.5.6), by uniformly randomly drawing individuals from the set of the n best individuals. In case of a deterministic local optimizer, the success of this approach depends on whether the best individuals have been altered since the candidates have last been subjected to local search. Otherwise the final local search results become the start solutions of the next search, which in general does not lead to any improvement. This LO-selection routine can be expected to be advantageous in which the search space sports a high degree of smoothness and convexity.

The second scheme is a variation of the elitarian operator in that it allows different pay-off partitions of the population to be considered. Since in this work we have used three partitions, the best, average and worst individuals, it is coined *tripartite* selection (ibid.). User-provided parameters n_b , n_a and n_w determine the selection sizes for the according partitions. Like before, this operator can be applied deterministically or probabilistically. The rationale behind this operator is that it increases the likelihood of different regions of search spaces to be exploited more intensely. Especially in the case of niching routines, there is no guarantee, however, to actually account for the nature of the search space such as its multimodality.

We have therefore considered a special case of the probabilistic elitarian operator, where the selection pool size is equal to the population. Under *egalitarian* selection each individual hence receives the same selection probability regardless of its fitness (*ibid.*).

In addition to the so-defined operators, also traditional selection operators such as tournament selection could be used. Distance-based operators, for instance, using similar heuristics as niching methods, may pose a promising alternative but have not been considered in this work (Hart, 1994; Ong and Keane, 2004).

5.3.3 Local Search Space Restrictions

As stated before, the local search operator is normally applied with a view to exploitation rather than exploration. Similarly to high mutation rates, a long-ranging local search can be associated with a high degree of schema disruption, thus strongly interfering with the mixing efforts of the GA. Moreover, except for well-natured⁹ functions, in which case a heuristic or hybrid search approach is generally not a good choice, long-reaching local searches can also be associated with increased computation times. For that reason, we have employed a restriction of the local search radius.

A straightforward approach is the introduction of fixed box constraints (Dürr, 2006b; Dürr et al., 2007a), for example, relative to the current parameter values of the candidate solutions to be locally modified. This approach depends on the ability of the employed local search method to handle parameter bounds. Otherwise, a surrogate constraint handling technique may have to be provided (see Section 4.4). For discrete optimizers, other search space reduction may also be qualified (see, for example, Krasnogor, 2002). In the presence of (linear or non-linear equality or inequality) constraints in the optimization problem, specific care has to be taken for their incorporation into the local search routine. We have not specifically considered such cases in this work, and instead relied on the same treatment of constraints as for the GA itself. That is, although the local search routine may be a preemptive constrained optimization technique, we have not formulated the problem accordingly but have employed the same surrogate formulations as in the GA.

In addition to a relative static restriction of the local search radius, dynamic, adaptive and self-adaptive radii may pose a viable alternative. For the *dynamic* case (Dürr, 2006b; Dürr et al., 2007b), for example, we have assumed an increase or decrease of the search radius in dependence of the generation index. Similarly as for the triggers (see Section 5.3.1), different signatures such as linear or exponential slopes may be appropriate.

When employing local search restrictions, statically or dynamically, care has to be taken that the employed radii do not impede the search in such a way that the resulted solution does not show any difference when coded back to discrete representation at all. Thus the search radius has to be chosen large enough to allow its defining region to cover other than the original solution with respect to the discrete correspondent. In the future, it may be useful to include additional information of the local search into the selection of radii. For example, if the current restriction selections routinely leads to solutions lying on the feasibility boundary constituted by the search space restriction, an according adaptive (yet moderate) increase of the radius may leverage the local search without strong schema-disruptive side-effects. Another alternative would be the introduction of a search space mapping similarly to decoders (see Section 4.4) such that the local search is conducted on a replacement problem that is chosen such to ensure a well-natured correspondence between local and global search.

5.4 Local Search Approaches

Obviously, the performance of the memetic algorithm is strongly impacted by its individual components, the global and the local search approaches and their interaction. Still, for the design chosen in this work, the influence of the local search approach can be considered a secondary operator because the broad idea is that the LO “helps” the genetic search, but the GA does not strongly depend on it. In that light, the proper choice of a local optimizer becomes a marginal issue, since any approach exhibiting

⁹For example, smooth, convex, analytic derivatives etc.

some form of hill-climbing behavior may be valuable. Three critical aspects are to be closely considered, however:

- (1) In order to be expedient, the memetic algorithm must not exceed the computation time of the stand-alone GA. Thus the number of function evaluation required for the local optimizer must be kept low. This is especially important in the regime of gradient-based methods, for which additional evaluation for the numerical determination of first and second-order information are required.
- (2) The ratio of accuracy and computation time should be in favor of lower computation times. Too high an accuracy can be considered wasteful. First, since it may not reflect the resolution achievable by the discretization used in the GA, and second because the solution obtained by the LO may be only an intermediate local optimum, quickly superseded by a another, globally better performing candidate. It is hence more promising to rather conduct multiple less computationally expensive local optimizations than to spend a high effort on one search step.
- (3) To limit the number of iterations in one local search run, it is favorable to employ approaches that allow for a hard limit of iterations.
- (4) In order for the proposed search radius limitations (see Section 5.3.3) to be applicable in a straightforward manner, a search approach supporting box constraints can be expected to pose an advantage.

In this work, we have examined the use of two (fundamentally) different local search approaches: (1) an algorithm based on an iterative applications of Newton's method on ad hoc quadratic replacement problems [sequential quadratic program (SQP)], (2) the heuristic downhill-simplex method proposed by Nelder and Mead (1965).

A brief review of both methods is provided in the following sections.

5.4.1 Sequential Quadratic Programming

The SQP approach is an algorithm for the numerical optimization of non-linear continuous problems that may be subject to both non-linear equality and inequality constraints. Wilson (1965) is often credited for the basic idea of SQP, though its popularity is due to the work by Han (1976, 1977) and Powell (1978a,b), and extensions by numerous other authors.¹⁰

The non-linear problem (NLP) can be stated as follows:

$$\begin{aligned}
 & \arg \min_{\mathbf{x}} && f(\mathbf{x}) \\
 & \text{subject to} && \mathbf{h}(\mathbf{x}) = \mathbf{0} \\
 & && \mathbf{g}(\mathbf{x}) \leq \mathbf{0} \\
 & \text{where} && f : \mathbb{R}^n \rightarrow \mathbb{R} \\
 & && \mathbf{h} : \mathbb{R}^n \rightarrow \mathbb{R}^m \\
 & && \mathbf{g} : \mathbb{R}^n \rightarrow \mathbb{R}^p,
 \end{aligned} \tag{5.6}$$

and where the functions f , \mathbf{g} and \mathbf{h} are assumed as belonging to C^2 , that is, they are twice continuously differentiable.¹¹ The basic concept of SQP can be described as an iterative procedure in which for each iterate a quadratic program (QP) is modeled about the iterate and solved. With an appropriate choice of the QP, SQP can be characterized as an extension of Newton's method to NLP.

“Perhaps the key to understanding the performance and theory of SQP is the fact that, with an appropriate choice of quadratic subproblem, the method can be viewed as the natural extension of Newton and quasi-Newton methods to the constrained optimization setting.” (Boggs and Tolle, 1995)

¹⁰There are a multitude of review papers and textbooks covering SQP, for example, Boggs and Tolle (1995), Gould and Toint (2000), Nocedal and Wright (2006, chap. 18) or Schittkowski and Yuan (2011).

¹¹Other smoothness conditions are sometimes applied. For their convergence considerations, Boggs and Tolle (1995), for example, assume the functions to be $\in C^3$.

Algorithm 5.1: Basic SQP Algorithm

```

1 State Set initial approximation  $(\mathbf{x}^{(0)}, \boldsymbol{\lambda}^{(0)}, \boldsymbol{\mu}^{(0)}, \mathbf{B}^{(0)})$ 
2  $k \leftarrow 0$ 
3 while not converged do
4    $\mathbf{d} \leftarrow$  Solution of Quadratic Problem
5   find  $\alpha \mid \phi(\mathbf{x}^{(k)} + \alpha \mathbf{d}) < \phi(\mathbf{x}^{(k)})$ 
6    $\mathbf{x}^{(k+1)} \leftarrow \mathbf{x}^{(k)} + \alpha \mathbf{d}$ 
7   Update  $\boldsymbol{\lambda}^{(k+1)}$  and  $\boldsymbol{\mu}^{(k+1)}$ 
8   Update  $\mathbf{B}^{(k+1)}$ 
9    $k \leftarrow k + 1$ 
10 return  $\mathbf{x}^* \leftarrow \mathbf{x}^k$ 

```

The flow of SQP is schematically shown in Algorithm 5.1 and can be summarized as follow: For each iteration step (k), a quadratic sub-problem is defined which approximates the original problem in the neighborhood of the current solution $\mathbf{x}^{(k)}$ (Algorithm 5.1, l. 4). The solution of this sub-problem is used to evaluate local properties of the search space, yielding the search direction \mathbf{d} for the subsequent step. The step length α is determined through a line search, applied to a merit function. In addition, both the equality and the inequality constraints are linearized about $\mathbf{x}^{(k)}$. The choice of the QP is critical as it determines the local convergence rate of the overall approach. It has the following general form:

$$\begin{aligned}
 & \arg \min_{\mathbf{d}} \quad \mathbf{r}^T \mathbf{d} + \frac{1}{2} \mathbf{d}^T \mathbf{B}^{(k)} \mathbf{d} \\
 & \text{subject to} \quad \nabla \mathbf{h}^T(\mathbf{x}^{(k)}) \mathbf{d} + \mathbf{h}(\mathbf{x}^{(k)}) = \mathbf{0} \\
 & \quad \quad \quad \nabla \mathbf{g}^T(\mathbf{x}^{(k)}) \mathbf{d} + \mathbf{g}(\mathbf{x}^{(k)}) \leq \mathbf{0},
 \end{aligned} \tag{5.7}$$

where \mathbf{d} is the descent vector, i.e., $\mathbf{d} = \mathbf{x}^{(k+1)} - \mathbf{x}^{(k)}$, and the vector \mathbf{r} and the $n \times n$ matrix $\mathbf{B}^{(k)}$ have to be chosen such that they closely resemble the neighborhood of $\mathbf{x}^{(k)}$, about which also the constraints are linearly approximated.

Instead of using a direct quadratic model of the objective function, an alternative QP, employing the (necessary) optimality conditions, is applied. The rationale for this approach is that it is less prone to inconsistencies under the linearization of non-linear constraints (Boggs and Tolle, 1995). Moreover, as discussed in the next subsection, using the quadratic optimality condition model, leads to a “natural” extension of Newton’s method, in which the Newton iterates also constitutes the iterates of SQP. The first-order necessary optimality conditions are conveniently posed, using Lagrangian multipliers, defined as follows:

$$\mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}) := f(\mathbf{x}) + \boldsymbol{\lambda}^T \mathbf{h}(\mathbf{x}) + \boldsymbol{\mu}^T \mathbf{g}(\mathbf{x}) \tag{5.8}$$

Additionally, in order for the necessary optimality conditions to apply to a local optimum \mathbf{x}^* , its constraints have to meet certain qualification conditions. Broadly speaking, these conditions ensure that the tangential behavior of the constraints is resembled by their linearized approximates (e.g., see Nocedal and Wright, 2006, pp. 315–320). One rather rigorous condition that can be applied, is the linear independence constraint qualification (LICQ), which considers the equality constraints and the active set of inequality constraints. Let the active set be defined as the index set of inequality constraints active for a candidate \mathbf{x}' :

Definition 5.1 (Active Set) *Given NLP (5.6) and a feasible point $\mathbf{x}' \in \mathbb{R}^n$. An inequality constraint g_i ($1 \leq i \leq p$) is said to be active at \mathbf{x}' if $g_i(\mathbf{x}') = 0$. The set of active inequality constraints*

$$\mathcal{I}(\mathbf{x}') := \{i \mid g_i(\mathbf{x}') = 0\} \tag{5.9}$$

is called active set of \mathbf{x}' .

Then,

Definition 5.2 (LICQ) *Given NLP (5.6), a point \mathbf{x}' and its active set $\mathcal{I}(\mathbf{x}')$. \mathbf{x}' satisfies the linear independence constraint qualification : \iff All elements of*

$$\{\nabla h_i \mid 1 \leq i \leq m\} \cup \{\nabla g_i \mid i \in \mathcal{I}(\mathbf{x}')\}$$

are linearly independent.

For a local minimum, we can hence define first-order necessary conditions, often termed Karush-Kuhn-Tucker (KKT) conditions, as follows:

Theorem 5.1 (First-order necessary conditions) *Given NLP (5.6) and a local minimizer $\mathbf{x}^* \in \mathbb{R}^n$ of NLP. If additionally LICQ (5.2) applies, then there exist unique vectors $\boldsymbol{\lambda}^*$ and $\boldsymbol{\mu}^*$ such that*

$$\nabla \mathcal{L}(\mathbf{x}^*, \boldsymbol{\lambda}^*, \boldsymbol{\mu}^*) = 0 \quad \text{stationarity} \quad (5.10a)$$

$$\begin{aligned} h_i(\mathbf{x}^*) &= 0 & (1 \leq i \leq m) \\ g_j(\mathbf{x}^*) &\leq 0 & (1 \leq j \leq p) \end{aligned} \quad \text{primal feasibility} \quad (5.10b)$$

$$\mu_j^* \geq 0 \quad \text{dual feasibility} \quad (5.10c)$$

$$\mu_j^* g_j(\mathbf{x}^*) = 0 \quad \text{complementary slackness} \quad (5.10d)$$

It should be noted that in general SQP does not depend on the feasibility of solutions candidates not even the start point (Boggs and Tolle, 1995). Thus the feasibility conditions [(5.10b), (5.10c)] are typically omitted.

Using the Lagrangian to model the QP subproblem (5.7), we have that

$$\begin{aligned} \arg \min_{\mathbf{d}_x} \quad & \nabla \mathcal{L}(\mathbf{x}^{(k)}, \boldsymbol{\lambda}^{(k)}, \boldsymbol{\mu}^{(k)})^T \mathbf{d}_x + \frac{1}{2} \mathbf{d}_x^T \mathbf{B}^{(k)} \mathbf{d}_x \\ \text{subject to} \quad & \nabla \mathbf{h}^T(\mathbf{x}^{(k)}) \mathbf{d}_x + \mathbf{h}(\mathbf{x}^{(k)}) = \mathbf{0} \\ & \nabla \mathbf{g}^T(\mathbf{x}^{(k)}) \mathbf{d}_x + \mathbf{g}(\mathbf{x}^{(k)}) \leq \mathbf{0}. \end{aligned} \quad (5.11)$$

We shall call this the KKT-QP. In addition to the descent direction \mathbf{d}_x , the solution of the QP yields the corresponding optimal Lagrangian multipliers, which—denoting them by $\boldsymbol{\lambda}^{(k+1)}$, $\boldsymbol{\mu}^{(k+1)}$ —constitute a possible choice for the subsequent SQP iteration (see Algorithm 5.1, l. 7). In order to exactly model the QP of the Lagrangian, the matrix $\mathbf{B}^{(k)}$ has to be chosen as the Hessian $\nabla^2 \mathcal{L}(\mathbf{x}^{(k)}, \boldsymbol{\lambda}^{(k)})$ (Algorithm 5.1, l. 8). This choice is also applied in the next paragraph, in which the local behavior of SQP is examined. Other choices, i.e., Hessian approximates, are discussed later.

Local properties of SQP

As mentioned before, the SQP approach can be broadly characterized as an iterative application of Newton's method to the optimality conditions of the NLP. This dualism poses a viable perspective for the local convergence properties of SQP, but also facilitates a number of practical considerations, for example, when choosing an appropriate Hessian approximate. The following derivation (cf. Nocedal and Wright, 2006, pp. 530–532) briefly motivates this view.

In order to examine the local properties of SQP, it is convenient to omit the inequality constraints of (5.6). This simplification is necessary to render the problem amenable to Newton's method. We shall

see at the end of this paragraph how to account for inequality constraints. Let the reduced problem be given as:

$$\begin{aligned} & \arg \min_{\mathbf{x}} && f(\mathbf{x}) \\ & \text{subject to} && \mathbf{h}(\mathbf{x}) = \mathbf{0} \\ & \text{where} && f : \mathbb{R}^n \rightarrow \mathbb{R} \\ & && \mathbf{h} : \mathbb{R}^n \rightarrow \mathbb{R}^m. \end{aligned} \quad (5.12)$$

The first-order KKT conditions can then be written as the following $n + m$ equations system, with $n + m$ unknowns presented by \mathbf{x} and $\boldsymbol{\lambda}$:

$$\begin{pmatrix} \nabla f(\mathbf{x}) - \nabla \mathbf{h}(\mathbf{x}) \boldsymbol{\lambda} \\ \mathbf{h}(\mathbf{x}) \end{pmatrix} = \mathbf{0}, \quad (5.13)$$

which by virtue of Theorem 5.1 is satisfied by solutions $\mathbf{x}^*, \boldsymbol{\lambda}^*$. A straightforward approach to find these solutions is the application of Newton's method, whose iterates of step k are posed as:

$$\begin{pmatrix} \nabla^2 \mathcal{L}(\mathbf{x}^{(k)}, \boldsymbol{\lambda}^{(k)}) & -\nabla \mathbf{h}(\mathbf{x}^{(k)}, \boldsymbol{\lambda}^{(k)}) \\ \nabla \mathbf{h}(\mathbf{x}^{(k)}, \boldsymbol{\lambda}^{(k)})^T & \mathbf{0} \end{pmatrix} \begin{pmatrix} \mathbf{d}_{\mathbf{x}} \\ \mathbf{d}_{\boldsymbol{\lambda}} \end{pmatrix} = \begin{pmatrix} -\nabla f(\mathbf{x}^{(k)}) + \nabla \mathbf{h}(\mathbf{x}^{(k)}) \boldsymbol{\lambda}^{(k)} \\ -\mathbf{h}(\mathbf{x}^{(k)}) \end{pmatrix} \quad (5.14)$$

In the proximity of $\mathbf{x}^*, \boldsymbol{\lambda}^*$, the Hessian $\nabla^2 \mathcal{L}^{(k)}$ can be assumed to be positive-definite.¹² In that case, the Newton step is well-defined and solvable. Thus the iteration update formula is:

$$\begin{pmatrix} \mathbf{x}^{(k+1)} \\ \boldsymbol{\lambda}^{(k+1)} \end{pmatrix} := \begin{pmatrix} \mathbf{x}^{(k)} \\ \boldsymbol{\lambda}^{(k)} \end{pmatrix} + \begin{pmatrix} \mathbf{d}_{\mathbf{x}} \\ \mathbf{d}_{\boldsymbol{\lambda}} \end{pmatrix} \quad (5.15)$$

On the other hand, if we consider the KKT-QP that is embedded in the SQP algorithm, which is modeled about iterate $\mathbf{x}^{(k)}$, we have that

$$\begin{aligned} & \arg \min_{\mathbf{d}} && \nabla f(\mathbf{x}^{(k)})^T \mathbf{d} + \frac{1}{2} \mathbf{d}^T \nabla^2 \mathcal{L}(\mathbf{x}^{(k)}, \boldsymbol{\lambda}^{(k)}) \mathbf{d} \\ & \text{subject to} && \nabla \mathbf{h}(\mathbf{x}^{(k)})^T \mathbf{d} + \mathbf{h}(\mathbf{x}^{(k)}) = \mathbf{0}, \end{aligned} \quad (5.16)$$

which under equality constraints is equivalent to (5.11). A unique solution to (5.16) is given, assuming that the constraint Jacobian has a full row rank and that the Hessian of the Lagrangian is positive-definite on the kernel of the constraint Jacobian. Specifically, for the solution $\mathbf{d}_{\mathbf{x}}, \mathbf{d}_{\boldsymbol{\lambda}}$, the following equation holds:

$$\begin{pmatrix} \nabla^2 \mathcal{L}(\mathbf{x}^{(k)}, \boldsymbol{\lambda}^{(k)}) \mathbf{d}_{\mathbf{x}} + \nabla f(\mathbf{x}^{(k)}) - \nabla \mathbf{h}(\mathbf{x}^{(k)}) \mathbf{d}_{\boldsymbol{\lambda}} \\ \nabla \mathbf{h}(\mathbf{x}^{(k)})^T \mathbf{d}_{\mathbf{x}} + \mathbf{h}(\mathbf{x}^{(k)}) \end{pmatrix} = \mathbf{0}, \quad (5.17)$$

In terms of iterates of Newton's method, the equations can hence be written as:

$$\begin{pmatrix} \nabla^2 \mathcal{L}(\mathbf{x}^{(k)}, \boldsymbol{\lambda}^{(k)}) & -\nabla \mathbf{h}(\mathbf{x}^{(k)}, \boldsymbol{\lambda}^{(k)}) \\ \nabla \mathbf{h}(\mathbf{x}^{(k)}, \boldsymbol{\lambda}^{(k)})^T & \mathbf{0} \end{pmatrix} \begin{pmatrix} \mathbf{d}_{\mathbf{x}} \\ \mathbf{d}_{\boldsymbol{\lambda}} \end{pmatrix} = \begin{pmatrix} -\nabla f(\mathbf{x}^{(k)}) \\ -\mathbf{h}(\mathbf{x}^{(k)}) \end{pmatrix}, \quad (5.18)$$

where $\boldsymbol{\lambda}^{(k+1)} = \boldsymbol{\lambda}^{(k)} + \mathbf{d}_{\boldsymbol{\lambda}}$. In other words, the iterates $\mathbf{x}^{(k+1)}, \boldsymbol{\lambda}^{(k+1)}$ can be obtained as one iterate of Newton's method applied to the NLP or as the solution to the quadratic program model of NLP. Thus using a local SQP variant, assuming initial solutions close to the final solutions and omitting line search, the convergence rate is determined by the one of Newton's method, which is quadratic.

The incorporation of the inequality constraints significantly increases the complexity of the QP subproblem. Particularly, Newton's method cannot be directly applied to the problem any longer. One

¹²Following from the second-order optimality condition (cf. p. 352, Nocedal and Wright, 2006).

approach to circumvent this problem is the introduction of slack variables \mathbf{z} , turning the inequality constraints \mathbf{g} into equality constraints,

$$\begin{aligned}\mathbf{g}(\mathbf{x}) + \mathbf{z} &= \mathbf{0} \\ \mathbf{z} &\geq \mathbf{0}.\end{aligned}$$

Because of the increased number of variables, active-set approaches are used more frequently. There, a working set is defined that consists of the equality constraints and an approximation of the active inequality constraints about the current iterate. Thus all constraints can be modeled as equality constraints rendering the so-modified QP subproblem again solvable with Newton's method. For each iterate, the working set has to be redetermined. Numerous techniques following this principle have been proposed such as gradient projection methods. "Warm-start" techniques are routinely applied in order to reuse computation results of previous iteration steps when applicable.

Hessian matrix approximate and update

Although using the exact Hessian of the Lagrangian in QP subproblem can be considered ideal in terms of the local convergence behavior, a number of problems in its computation can be attributed to it. Most importantly: (1) Especially when considering points not close to the local optimum, the Hessian may not be positive-definite, thus averting the (efficient) solution of the QP. (2) In many contexts a straightforward determination of the Hessian is either computationally expensive or even impossible. (3) Depending on the steplength by which a new candidate is obtained, the local properties may not change drastically, so that a full recomputation of the Hessian is not required. These considerations lead to an update scheme in which the Hessian approximate for the current iteration step is based on the former Hessian iterate, for example, using a Broyden-Fletcher-Goldfarb-Shanno (BFGS) approach (Broyden, 1967; Fletcher and Powell, 1963; Goldfarb, 1970; Shanno, 1970). The Hessian or its approximate may be postconditions in order for it to become positive-definite. Another approach is to apply a modified (damped) BFGS approach, which ensures the positive-definiteness, provided the initial Hessian is positive definite (Powell, 1978a). This method, exhibits a R-superlinear convergence rate; a formal convergence proof has not been presented. Numerous alternative update scheme have been proposed (e.g., Han, 1976; Powell and Toint, 1979; Nocedal, 1980; Tapia, 1988; Conn et al., 1991; Fletcher, 2006) and successfully applied to line-search or trust-region SQP variants (briefly explained later) in which due to the locality constraint, the positive-definiteness requirement is relaxed.

An alternative to approximating the full Hessian is a reduced Hessian approximation scheme (Byrd and Nocedal, 1990). This can be motivated by acknowledging the fact that the positive definiteness of the Hessian is required only on the kernel of the Jacobian of the constraints. Broadly speaking, the reduced schemes are aimed at decoupling the computations of the QP multipliers and the descent vector. That way the contributions of the Hessian on the kernel of the constraint matrix can be dealt with separately. In addition to a reduced computation time, the probability of yielding well-conditioned approximates drastically increases in comparison to a full approximation, and a two-step superlinear convergence rate can be achieved. Its application to a QP including inequality constraints is less straightforward, however.

Line search and merit function

The former considerations are concerned with the *local convergence* properties of SQP. That is, given a candidate sufficiently close to the solution, the SQP can be proved to converge superlinearly (Han, 1976). In order to render it practically useful, a *global convergence* safeguard has to be incorporated, which requires an additional effort and which can be expected to reduce the convergence rate. A common approach is the use of a line search in combination with a merit function. The solution of the QP subproblem \mathbf{d}_x can then be considered as the descent direction. Instead of updating the iterate with

a unit length toward this direction, an appropriate steplength is determined for which a merit function yields a sufficient decrease. Let ϕ be this merit function and α the step length parameter to be sought:

$$\begin{aligned} \mathbf{x}^{(k+1)} &:= \mathbf{x}^{(k)} + \alpha \mathbf{d}_x, \text{ where} \\ \phi(\mathbf{x}^{(k+1)}) &<_s \phi(\mathbf{x}^{(k)}). \end{aligned} \quad (5.19)$$

Similarly to related approaches discussed earlier (see Section 4.4), the merit function ϕ consists of the objective and penalty contributions due to constraint violations. A popular choice is the following ℓ_1 form:

$$\phi(\mathbf{x}) := f(\mathbf{x}) + \rho (|\mathbf{h}(\mathbf{x})|_1 + |\overline{\max}(\mathbf{g}(\mathbf{x}), \mathbf{0})|_1),$$

where ρ is the penalty parameters for the equality and inequality constraints, respectively. The operator $\overline{\max}$ is defined as a component-wise maximum.¹³ A line search step is considered sufficient—and hence accepted—if and only if the following condition holds:

$$\phi(\mathbf{x}^{(k+1)}) \leq \phi(\mathbf{x}^{(k)}) + \eta \alpha D_{\mathbf{d}_x}(\phi(\mathbf{x}^{(k)})), \quad (5.20)$$

where $D_{\mathbf{d}_x}$ is the directional derivative¹⁴ and $\eta \in [0, 1/2]$ is a user-provided parameter. Like in the application of penalty methods to evolutionary algorithms, the selection of appropriate penalty parameters is critical and should, for example, present a scaling that leads to a higher penalty of currently active constraints. A multitude of alternative merit and penalty functions exist—for example, see overviews given by Boggs and Tolle (1995) or Nocedal and Wright (2006, chap. 17 and 18). Although generally neither of the merit functions converge to the local optimum, the global convergence of algorithms employing them can be shown. Broadly, in order for the line search to be considered efficient, it should conserve the superlinear or two-step superlinear properties of SQP. But even close to the solution, unit steplength (α) may not be obtained under the sufficient decrease condition (5.20), potentially leading to a significant degradation of the convergence rate.¹⁵ To recover superlinear convergence, second-level corrections may be applied.

Trust-region approach

As an alternative to the line search approach, a trust-region method can be used instead (Vardi, 1985; Celis et al., 1985; Byrd and Schnabel, 1987). There, the locality of the QP solution is ensured by adding extra (box) constraints to the KKT-QP. In order to render the problem well-conditioned, the trust region constraint and the original constraints must be consistent. Relaxation methods are hence commonly introduced. In addition to an update strategy to adapt the trust-region constraints to a new iterate, a method for the acceptance or the rejection of QP candidates has to be defined. Because of the strict locality that can be enforced on QP candidate solutions, this method is especially well suited in conjunction with partial Hesse approximations.

Implementation

In this work, we have employed two alternative implementations of SQP, for which details can be found in (Spellucci, 1998a,b) and (MathWorks, 2004, chap. 3) respectively. Both implementations use a line-search approach, using slightly different merit functions. The QP is solved in both variants using (different) active-set methods, employing different Hessian approximation techniques.

In the present version of the proposed memetic algorithm, only box constraints, in order to limit the local optimization radius, are employed. Other problem-specific constraints are dealt with in the same fashion as for the GA, using a merit function approach. To exploit the full potential of SQP, a dual,

¹³I.e., given $\mathbf{x}^1, \mathbf{x}^2, \dots, \mathbf{x}^l \in \mathbb{R}^m$,

$$\mathbf{y} = \overline{\max}(\mathbf{x}^1, \mathbf{x}^2, \dots, \mathbf{x}^l) : \iff y_i = \max\{x_i^1, x_i^2, \dots, x_i^l\} \quad (i = 1, \dots, m).$$

¹⁴which exists, although ϕ is not differentiable.

¹⁵This is often referred to as the Maratos effect (Maratos, 1978).

optimizer-tailored representation of the constraints may be more appropriate. This would, however, pose a further representation alignment problem, and has hence not been considered in this work.

5.4.2 Nelder-Mead Down-hill Simplex Method

Based on a related approach by Spendley et al. (1962), the down-hill simplex method was first introduced by Nelder and Mead (1965). It is an unconstrained, derivative-free, heuristic optimization technique, which has been successfully applied to a large number of problems (for illustration, see the list of applications in chemometrics compiled by Deming and Morgan, 1983). The convergence rate and hence the number of function evaluations can be generally expected to be larger compared to gradient-exploiting methods such as the SQP approach. However, even with a few iteration steps, the Nelder-Mead (NM) algorithm is often able to significantly improve the start value. As a disadvantage, however, this method is not guaranteed to actually converge to a critical point (McKinnon, 1998).

The NM method is an iterative procedure, which employs the notion of a simplex, the simplest possible polytope of the given search space. A simplex is therefore convex but not necessarily regular. To illustrate: In a one-dimensional space, a simplex is defined by a line, in two dimensions by a triangle. Generally, for an n -dimensional search space, we obtain the following $n + 1$ vertices:

$$\{x_i\}_{i=0}^n.$$

With each iteration step, the vertices of the simplex are modified such that it moves towards and eventually encompasses locally optimum regions of the search space. This behavior is modeled through simple geometric operations, called *reflection*, *contraction*, *expansion* and *shrink*. The algorithm proceeds until a pre-defined termination criterion is met such as a maximum number of iterations, a sufficient improvement of the merit function or other problem-tailored criteria.

Given the following unconstrained optimization problem:

$$\begin{aligned} \arg \min_{\mathbf{x}} \quad & f(\mathbf{x}) \\ \text{where} \quad & f : \mathbb{R}^n \rightarrow \mathbb{R}. \end{aligned} \tag{5.21}$$

The NM method does not account for constraints. Thus similarly as in evolutionary algorithms, constrained problems require the inclusion of according measures (see Section 4.4). In contrast, to EAs, neither parameter bounds are dealt with.

The Nelder-Mead approach is formally summarized in Algorithm 5.2. The details of the individual steps are as follows: In the first step, an initial simplex is generated. As a local search method, the NM algorithm depends on the definition of a start value. Therefore, a common initial simplex is obtained from a linear combination of the start vector and the (scaled) unit vectors of the search space:

$$\mathbf{x}_i := \mathbf{x}_0 + \lambda \mathbf{e}_i \quad (i = 1, \dots, n), \tag{5.22}$$

where \mathbf{e}_i is the i -th unit vector, $\lambda > 0$ is a scaling factor for the initial size of the simplex, and \mathbf{x}_0 is the start vector. The vertices of the simplex are then ordered according to their objective function values:

$$f(\mathbf{x}_{i_k}) < f(\mathbf{x}_{i_l}) \rightarrow i_k < i_l \tag{5.23}$$

The first simplex operation is called *reflection*. For that, the centroid $\bar{\mathbf{x}}$ of the convex hull of all but the maximum (worst) points is evaluated,

$$\bar{\mathbf{x}} := 1/n \sum_{i=0}^{n-1} \mathbf{x}_i$$

The worst point (\mathbf{x}_n) is reflected along the line through the centroid (see Figure 5.6a):

$$\mathbf{x}_r := (1 + \alpha)\bar{\mathbf{x}} - \alpha\mathbf{x}_n, \tag{Reflection}$$

Algorithm 5.2: Nelder-Mead Down-hill Simplex

```

1 initialize simplex:  $\mathbf{x}_i \leftarrow \mathbf{x}_0 + \lambda \mathbf{e}_i$   ( $i = 1, \dots, n$ )          /* where  $\mathbf{e}_i$  is the  $i$ -th */
                                     /* unit vector of the parameter space */
2 while not termination criterion met do
3   sort vertices according to function values          /* see (5.23) */
4    $\mathbf{x}_r \leftarrow \text{reflection}$           /* see (Reflection) and Figure 5.6a */
5   if  $f(\mathbf{x}_0) \leq f(\mathbf{x}_r) < f(\mathbf{x}_{n-1})$  then
6      $\mathbf{x}_n \leftarrow \mathbf{x}_r$           /* set reflected point as new maximum (worst) */
7   else if  $f(\mathbf{x}_r) < f(\mathbf{x}_0)$  then
8      $\mathbf{x}_e \leftarrow \text{expansion}$           /* see (Expansion) and Figure 5.6b */
9     if  $f(\mathbf{x}_e) < f(\mathbf{x}_r)$  then
10       $\mathbf{x}_n \leftarrow \mathbf{x}_e$           /* set expanded point as new maximum (worst) */
11    else
12       $\mathbf{x}_0 \leftarrow \mathbf{x}_r$           /* set reflected point as new maximum (worst) */
13  else
14     $\mathbf{x}_a \leftarrow \text{contraction}$           /*  $\neg$  (lines 5  $\vee$  7)  $\rightarrow f(\mathbf{x}_r) \geq f(\mathbf{x}_{n-1})$  */
15    if  $f(\mathbf{x}_a) < f(\mathbf{x}_n)$  then
16       $\mathbf{x}_n \leftarrow \mathbf{x}_a$ 
17    else
18       $\mathbf{x}_i \leftarrow \text{shrink}$   ( $i = 2, \dots, n$ )

```

where $\alpha > 0$ is the reflection constant. It defines the distance of the reflected point relative to the distance between the centroid and the maximum point. A common choice is a value of 1.

If the function value of the newly computed reflected point \mathbf{x}_r lies between the minimum and the maximum point, i.e., if

$$f(\mathbf{x}_0) \leq f(\mathbf{x}_r) < f(\mathbf{x}_{n-1}),$$

\mathbf{x}_r becomes the new maximum point, thus flipping and changing the direction of the simplex. A new iteration step is then performed.

If otherwise the reflected point forms a new minimum, a new promising search direction has been discovered. In that case, the reflected point is expanded farther in order to find a point with an even lower function value:

$$\mathbf{x}_e := \gamma \mathbf{x}_r + (1 - \gamma) \bar{\mathbf{x}}, \quad (\text{Expansion})$$

where $\gamma > 1$ defines the expansion length as a ratio relative to the distance between the centroid and the reflected point. It defaults to 2. The *expansion* operation is illustrated in Figure 5.6b.

If the result of the expansion is indeed a new minimum point, \mathbf{x}_e replaces the old minimum \mathbf{x}_0 . This variant of expansion, which is also proposed in the original work of Nelder and Mead, is sometimes referred to as *greedy expansion* because it enlarges the simplex even if the former reflected point outperforms the expanded one. Although this procedure offers a robustness advantage on highly discontinuous functions, it has a negative impact on the convergence time, for example, due to degenerated simplices¹⁶, to which this may lead. In modern NM implementations, including the one implemented for this work, the minimum point \mathbf{x}_0 is replaced with the expanded point only if the expanded point is smaller than the reflected one. Otherwise, \mathbf{x}_r , the reflected point, is introduced as the new minimum point. In both variants, if the expansion did not yield a point with lower function values than \mathbf{x}_0 , the reflected point \mathbf{x}_r is selected as new minimum vertex. After the expansion, a new iteration step is started.

As indicated before, expansion is not performed if the reflection step led to a point that exhibits a larger or equal function value compared to the second largest vertex \mathbf{x}_{n-1} . That is to say that either the

¹⁶in the sense that the simplex becomes arbitrarily narrow, for example

reflection direction is not towards a valley or that the simplex already encompasses a local minimum. A *contraction* operation is applied to resolve this situation. If the reflected point has a smaller value than the maximum vertex

$$f(\mathbf{x}_r) < f(\mathbf{x}_n), \quad (5.24)$$

which could be thought of as too far a reflection, then an *outside contraction* is performed, as shown in Figure 5.6c:

$$\mathbf{x}_a := \beta \mathbf{x}_r + (1 - \beta) \bar{\mathbf{x}}, \quad (\text{Outside Contraction})$$

where $0 > \beta > 1$ is the contraction constant, with a default value of $1/2$. If the resulted contraction point \mathbf{x}_a is equal to or outperforms the reflected point, it replaces the maximum point \mathbf{x}_n , starting a new iteration step.

If, on the other hand, not even the reflected point resulted in an improvement,

$$f(\mathbf{x}_r) \geq f(\mathbf{x}_n), \quad (5.25)$$

a contraction inside the original simplex (*inside contraction*) is performed (see Figure 5.6d):

$$\mathbf{x}_a := \beta \mathbf{x}_n + (1 - \beta) \bar{\mathbf{x}}. \quad (\text{Inside Contraction})$$

If the obtained point \mathbf{x}_a possesses a lower function value than the current maximum vertex \mathbf{x}_n , it replaces \mathbf{x}_n , finalizing the iteration step.

If contraction fails, that is, if either outside contraction did not yield a point with a smaller merit value than the reflected one, or if inside contraction did not result in a point with a smaller function value than the currently worst vertex, a shrink step is performed. With this step, all but the minimum vertex are translated along their edge with the minimum point, decreasing the volume of the simplex.

$$\mathbf{x}_i := \mathbf{x}_0 + \sigma(\mathbf{x}_i - \mathbf{x}_0) \quad (i = 1, \dots, n) \quad (\text{Shrink})$$

According to Nelder and Mead (1965), failed contractions can be associated with situations in which a deformed¹⁷ simplex is located in a curved valley. In that case, additional contraction steps are not suited to adequately adapt the simplex. As already pointed out by Nelder and Mead (1965), a failed contraction can be considered a rare occasion. This conjecture was confirmed by Torczon (1989) through a large number of numerical experiments. Later, Lagarias et al. (1998) formally proved that under strictly convex functions, no shrink step occurs.

In direct search, the formulation of a robust termination criterion that does not prevent the convergence to a stationary point is far more critical compared, for example, to gradient-based optimization approaches like SQP. In their original work, Nelder and Mead (1965) proposed to halt the algorithm as soon as the standard error of the vertices falls below a user-defined threshold (ϵ):

$$\frac{1}{n} \sum_{i=0}^n (f(\mathbf{x}_i) - \bar{f})^2 < \epsilon,$$

where \bar{f} denotes the average objective value due to all vertices. Alternative criteria, based on movement of the simplex (Parkinson and Hutchinson, 1972) or its relative size (Woods, 1985), have been proposed. Combinations of termination criteria are sometimes employed, often supplemented with safeguard such as a maximum number of iterations.

Despite its often good performance acknowledged by numerous practitioners, little attempt has been made to formally analyze its convergence behavior. It has been found that even for strictly convex smooth functions (with two variables), the NM method is not guaranteed to converge to the local minimizer but may stagnate on non-stationary points (Torczon, 1989; McKinnon, 1998). Lagarias et al. (1998) proved that in one dimension and for strictly convex function, the NM method actually converges to the minimizer. In the two-dimensional case, however, only a convergence to a zero-diameter

¹⁷We use the term *deformed* here to distinguish it from *degenerated* simplices, which, provided the start simplex is not degenerated Lagarias et al. (1998).




Figure 5.6: Nelder-Mead down-hill simplex operations.

simplex was shown. Lagarias et al. (1998) and McKinnon (1998) conjecture that this behavior can be attributed to the fact that the NM is design such that simplices adapt to the local landscape. The different convergence results for related approaches, not employing such a simplex adaptation (Torczon, 1997), support this view.¹⁸ The convergence rate of NM is in the order of m -step linear. Very recently, Lagarias et al. (2012) proved that for a restricted NM, without expansion steps, convergence to the local minimizer is given if the Hessian of the search function is positive-definite everywhere and the initial simplex is non-degenerated. This result, which puts harsh conditions on the functions under considerations, can be interpreted as a local-NM convergence proof, since no expansions can be expected in the vicinity of optimizers. The authors conjecture that the experiments of McKinnon (1998) fail because of the missing ubiquitous positive-definiteness of their Hessians.

A large number of variants to the original Nelder-Mead approach have been suggested (Parkinson and Hutchinson, 1972). Especially, the shrink operations has been a subject of numerous researches and improvement proposals. Another point of interest is the shape of the simplex, which especially under noisy conditions may be deformed after a number of iterations. To circumvent this issue, numerous restart methodology have been proposed (Kolda et al., 2003; Humphrey and Wilson, 2000; Kelley, 1999; Parkinson and Hutchinson, 1972). There, the simplex is reinitialized, thus curing the deformation.

As mentioned before, the NM method is an unconstrained approach. In order to handle constraints, penalty methods have been proposed. Box (1965), however, reported that such an approach is prone to a simplex deformation and as a consequence to a convergence to a non-stationary point.¹⁹ Subrahmanyam (1989) aimed at improving the simplex stability under penalty methods by introducing a delayed reflection. Their approach, however, may lead to a significant increase of function evaluations. A premature convergence to non-stationary points is often associated with deformed simplices. To counteract this issue, many authors have devised NM variants in which a restart scheme is applied; that is, if the simplex becomes close to degenerated, a new initial simplex is constructed, and the NM procedure is freshly restarted. One of the difficulties with this approach is the discovery of a deformed simplex. Kelley (1999) uses a stagnation indicator based on a sufficient decrease metric. If it fails, a size-reduced simplex, oriented in descent direction is formed and the algorithm is restarted. This approach has been demonstrated to minimize the test functions by McKinnon (1998). In the work of Nazareth and Tseng (2002), NM is modified such that its configuration parameters follow the golden-section rule. Inspired by golden-section search, this approach allows for a straightforward convergence proof for a wider range of one-dimensional functions than in the related efforts, for example Lagarias et al. (1998). In addition, the authors slightly modify the criteria for the acceptance of the individual steps (i.e., reflection, expansion and contraction) by incorporating safeguards. The resulted descent is called *fortified descent* by the authors.

A comprehensive study on the sensitivity of the configuration parameters and on variations of the NM study was conducted by Parkinson and Hutchinson (1972). There, improvements were specifically aimed at regularizing the simplex in order to reduce the number of function evaluations. The dimensionality of the search space has been a major concern for applications of the NM method, an issue that was recently confirmed by Han and Neumann (2006).

5.5 Benchmark Problems

In order to assess the performance of the proposed MA and to identify adequate configuration parameters, several experiments on established benchmark functions have been conducted (Dürr, 2006b; Dürr et al., 2007b,a), some of which are summarized in the following. The functions are chosen such that they exhibit a certain “hardness” for EAs in that they possess a high level of modality or epistasis. As indicated before, the MA proposed here can be regarded as an endeavor to improve the EA in terms of convergence speed and robustness. To verify the success in achieving this goal, we have consequently compared the convergence behavior of the (local tournament) GA and the MA. The first experiment

¹⁸Also note the similar convergence results obtained for an NM variant by Woods (1985).

¹⁹He hence proposed an alternative method with termed *complex* method, in which convex feasibility regions are assumed (ibid.).

was conducted on a (maximization) variant of Shekel's Foxholes problem (see Appendix F) on a two-dimensional decision space. The according function exhibits a number of steep local optima, a situation that poses a significant difficulty for many optimizers. The population sizes of both the GA and the MA were set to 41, the maximum number of generation was 150. A crossover probability of 0.6 and, for the sake of an increased reproducibility, no mutation were used for the binary-coded GA. A coarse coding precision of 0.1 (i.e., one decimal place) was employed. The following MA configuration parameters are used: (1) SQP as the local optimization approach. (2) A *static* trigger, with an interval of 5 generations and an initial waiting period of 10 generations. (3) A *tripartite* selection with probabilities of 0.05 within the best, the average and the worst performing population partitions. (4) A constant optimization radius of 0.5 in each parameter.




Figure 5.7: Comparison of performances of GA and MA on two-dimensional Shekel's Foxholes function over 100 independent runs: the shaded regions show the range of minimum and maximum best fitness values of all runs. The thick lines denote the average. For the MA, a static trigger and a tripartite selection scheme were used.

Figure 5.7 shows the offline performances per generation of the MA and the GA after 100 independent runs. The shaded regions depict the range of the best performing individual out of these runs, i.e., the best and the worst performing champion candidates. The thick lines in between displays the mean fitness. To adequately compare the runtime behavior and also account for a potential overhead, for example, due to the increased communication demands, the actual computation time of both approaches is used as a performance measure. The abysmal performance of both the GA and the MA is due to the laborious interaction of the fitness evaluation and the runtime overhead associated with it: For these first tests, the SQP implementation of the Matlab Optimization toolkit (MathWorks, 2004) was employed. Therefore also the fitness evaluation is conducted using Matlab. The genetic algorithm, however, uses Python as its interfacing language; passing of parameters and results is performed through a simplistic remote call interface, leading to a runtime penalty of a factor of about 100. Although not appropriate for real-world applications, this property does not impact the relative performance measure. As shown in the graph, both the GA and the MA exhibit a high, albeit not a one-hundred-percent success rate in maximizing the function ($f(\mathbf{x}^*) \approx 25$). The MA shows a slightly improved average convergence rate and a smaller maximum deviation.

The second test was conducted on a variant of Rastrigin's problem. Particularly, the 20-dimensional function as presented in Appendix F was to be maximized. This function exhibits a strong multimodality but a negligible degree of epistasis. Due to their pronounced gradients, the local optima can be expected to constitute strong attractors for local optimizers. The population sizes for the GA and the MA were

set to 60. The number of generations was limited to 150. The following MA configuration was applied: (1) A *dynamic* asymptotic trigger ($\alpha : 0.5$), with an initial probability of 0.05 at generation 10 and a final probability of 0.3. (2) An *egalitarian* selection with a probability of 0.1. (3) And a constant optimization radius of 0.5 in each parameter.

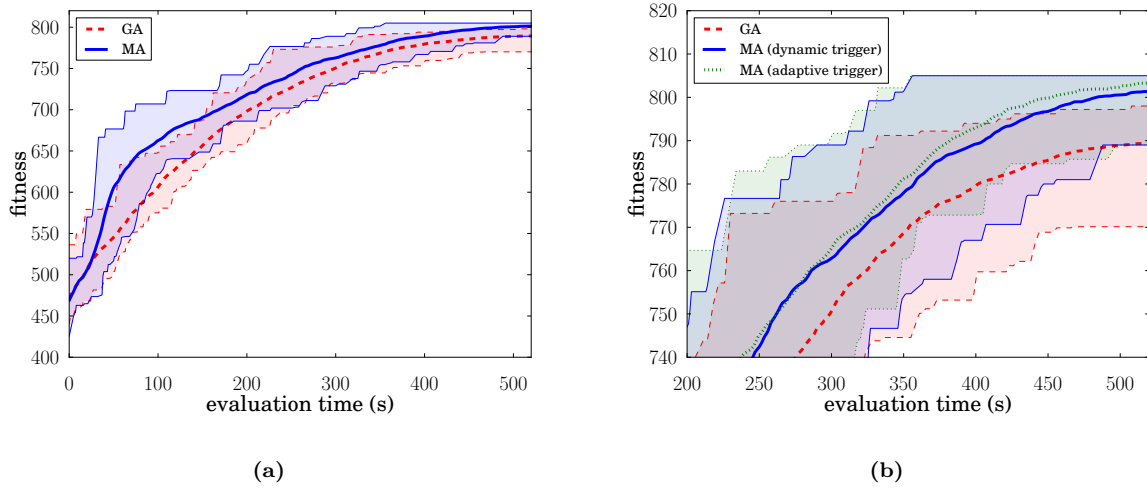


Figure 5.8: Comparison of performances of GA and MA on 20-dimensional Rastrigin function over 50 independent runs. (a) Maximum, mean and minimum best fitness comparison of GA and MA using a dynamic asymptotic trigger and an egalitarian selection. (b) Comparison of final stage of GA and MA as in (a) with MA using an adaptive trigger.

Figure 5.8(a) shows the comparison of 50 independent GA and MA runs. As before, the shaded regions demonstrate the range of maximum and minimum best fitness values; the thick lines denote the mean best fitness values. The MA outperforms the GA in terms of convergence rate standard and maximum deviation. In none of the GA runs the global optimum of about 807 was found, whereas the MA showed a success rate close to one.

In a second test performed on the same problem, using the same selection procedure and optimization radii, an *adaptive* trigger was used. There, the local optimization operator is executed if the fitness improvement has dropped to values below three percent compared to the ten previous generations and an additional repose interval of five generations has elapsed. Figure 5.8(b) presents a comparison of the dynamic and the adaptive trigger and the GA for the mid to final stages of the runtime of the algorithms. Both MA configurations significantly outperform the GA. In this example, the adaptive trigger provides a minor advantage over the dynamic one since it further reduces the variance of best fitness values. It additionally shows a slightly decreased maximum deviation throughout the entire runtime.

Another series of tests was conducted on Griewank's function, which shows a similarly multimodal structure as the previous test on Rastrigin's function but also introduces a moderate degree of epistasis (i.e., a non-separability of parameters). This property can be expected to render the problem more difficult for the genetic algorithm (see Section 4.10). In the first test, a ten-dimensional Griewank function, a population size of 200, a maximum generation number of 100 and the following MA settings were used: (1) A *dynamic* linear trigger, with an initial probability of 0.1 at generation 10 and a final probability of 0.25 at generation 65 (i.e., $\alpha : 1/65$). (2) An *egalitarian* selection with a probability of 0.15. (3) And a constant optimization radius of 1 in each parameter.

As depicted in Figure 5.9, the behavior of the GA and the MA are at first very similar, with a slightly steeper descent of the GA. In a later stage, in between 900 and 1000 seconds, the MA surpasses the GA, initially showing a larger maximum deviation. After about 1900 evaluation seconds, the MA achieves

values very close to global optimum of 0. Additionally, the variance and the maximum deviation are drastically reduced compared to the GA.

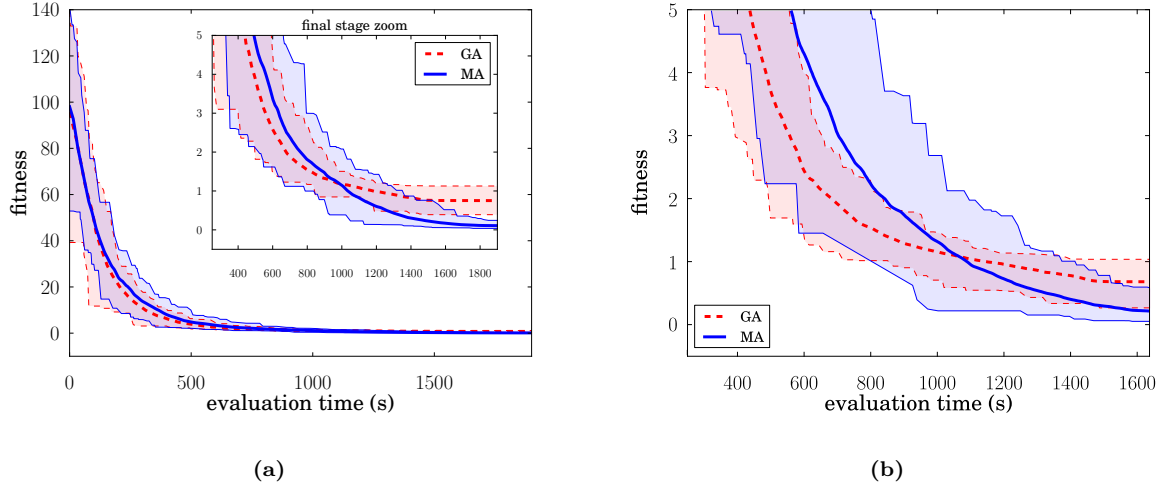


Figure 5.9: Comparison of performances of GA and MA on ten-dimensional Griewank function over 50 independent runs. (a) Maximum, mean and minimum best fitness comparison of GA and MA using a dynamic linear trigger and an egalitarian selection; upper right inset provides a close-up view of the final stage. (b) Final stage for the same test problem, albeit using an adaptive trigger for the MA.

The same test was repeated with an *adaptive* trigger: A local search step was executed whenever the fitness improvement within an interval of three generations stayed below five percent and at least four generations without local search have passed. As demonstrated in Figure 5.9, the MA shows a slightly worse behavior compared with the tests using the dynamic trigger in that it outperforms the GA only after 1100 seconds and in evincing both a higher variance and maximum deviation compared with the dynamic trigger. This performance is in contrast to the former maximization test on Rastrigin's function, where the adaptive trigger provided a minor advantage over the dynamic one. This observation suggests that under certain conditions (e.g., extreme multimodality in conjunction with epistasis), the periodic application of local search operations may become beneficial even if the GA alone steadily progresses. In other situations, local search is best applied whenever the GA stagnates.

In a last test on Griewank's function, the Nelder-Mead down-hill simplex technique was used instead of the SQP method of the previous tests. Since the NM method does not natively handle constraints, the following pull-back mechanism is used: If, when checking the termination criterion (see Section 5.4.2) the obtained point violates the box constraints given through the problem or the local optimization radius, the respective parameters are pulled back to the feasibility boundary and the algorithm is terminated. This simplistic is aimed at guaranteeing the feasibility of solutions while at the same time circumventing the problems associated with the application of penalty methods to the NM algorithm. In general, the so-obtained point is no local optimum. A more sophisticated line-search could be incorporated to ensure the local optimality, which for the sake of saving computation time we abandoned in the scope of this work. The dimensionality of the Griewank function was increased to 20 for this test. Consequently, to account for the aggravated complexity, also the maximum number of generations was increased, to 1000. For the MA configuration, we used: (1) The Nelder-Mead method as the local optimization approach. (2) A *static* trigger, with an interval of 20 generations and an initial waiting period of 40 generations. (3) An *egalitarian* selection with a probability of 0.075. (4) A constant optimization radius of 0.5 in each parameter.

The convergence graph comparing the MA and GA on ten independent runs is shown in Figure 5.9. Since both the fitness evaluations and the local optimizations were performed in parallel, the number of function evaluations instead of the computation time are compared, which for the purpose of com-

paring the MA and the GA performance constitutes an equivalent metric as the previous real time measurement. The MA considerably outperforms the GA, which can be assumed to suffer from the combination of high multimodality, high dimensionality and the epistasis presented by the function. After about 175 000 function evaluations, the MA achieves on average champion function values on the order of 10^{-6} , whereas the corresponding fitness values of the GA drops only marginally below one. Moreover, the MA exhibits a steady descent behavior, which can presumably be attributed to a steady discovery of local optima or points close to those.




Figure 5.10: Comparison of performances of GA and MA, using the Nelder-Mead method on twenty-dimensional Griewank function over ten independent runs. Maximum, mean and minimum best fitness values in logarithmic scales. Because fitness evaluations are performed in parallel, the number of evaluations instead of the computation times are compared.

In another benchmark example, the unimodal Rosenbrock function in a 30-dimensional decision space was minimized. As shown in Appendix F, this function features a relatively flat plateau in the vicinity of its minimizer **1**. Additionally, the minimum is located in a trench. Both properties pose a difficulty for common optimizers, which aggravates with an increasing input dimension. A population size of 100 and a maximum generation count of 900 were used in this example. An SQP and the following configuration were applied for the MA: (1) A *static* trigger, with an interval of 10 generations and an initial waiting period of 20 generations. (2) An *egalitarian* selection with a probability of 0.1. (3) And a constant optimization radius of 6 in each parameter.

The convergence graph for 25 independent runs is presented in Figure 5.11, like before showing the offline fitness ranges and the mean fitness per generation. Since fitness evaluations are conducted in parallel, the runtime is denoted as number of function evaluations. The MA shows both a significantly improved convergence and a reduced variance, with solutions close to, albeit not exactly at the global optimum.

For the last example, the 20-dimensional Ackley function was minimized Appendix F. This problem has a global minimum of 0 at **0**. It consists of a large number of global optima, which pose a particular difficulty at the outer regions of the search space, where the function does not descent toward the global optimum. In addition to a population size of 100 and a maximum generation number of 700, the following MA parameters were used: (1) A *dynamic* trigger, with an initial probability of 0.05 at generation 30 and asymptotic progression with a maximum of 0.14 reached at generation 500. (2) An *egalitarian* selection with a probability of 0.08. (3) A constant optimization radius of 6 in each parameter.

Figure 5.12, which shows the analysis of 25 independent runs, indicates a slight performance advantage of the MA over the GA. It is interesting to note that in the mid-term phase (about 10,000 to




Figure 5.11: Comparison of performances of GA and MA on 30-dimensional Rosenbrock function over 25 independent runs: As before, the shaded regions show the range of minimum and maximum best fitness values of all runs. The thick lines denote the average. The fitness value is depicted in logarithmic scale.

20,000 function evaluations), the GA outperforms the MA on average. This behavior can be supposedly attributed to the fact the local optimizer vastly descends into local optimum, slightly impairing the exploratory capacity of the GA. This disadvantage, however, is outweighed in the later stage in which the success rate of the MA becomes higher than that of the GA, and in two runs the global minimum was obtained.

The tests summarized here and the additional tests presented elsewhere (Dürr, 2006b; Dürr et al., 2007b,a), show the potential of the MA in improving the convergence rate and the precision of the genetic algorithm. The chosen benchmark problems pose certain difficulties to the GA including a high dimensionality, multimodality and epistasis. Similar properties can be expected to also apply to the optimization of lithographic process conditions, albeit the tests do not allow for a direct deduction of the actual “real-world” performance. In addition to verifying the performance improvement gained by the hybridization, another goal of the tests was the determination of an adequate configuration of the MA. Particularly, different trigger, selection and optimization radius settings were probed for. Only slight performance differences for the different trigger functions were observed, with a marginal advantage of the *dynamic* and the *adaptive* trigger over the *constant* version. This finding is in line with the conjecture that early optimization stages should emphasize exploration whereas a later phase profits from an intensified search. In order for the GA to take advantage of the local improvements, it has to be granted a sufficient number of generations in which it mixes the local modifications with the remaining candidates. It is thus important to trigger the local search operation with a reasonable frequency. Dürr (2006b), for example, conjectures that an application of local search to more than a fourth of the total generations drastically impairs the exploration power of the GA. Judging from the benchmark tests, an interval of ten to twenty generations, depending on the population size, appears to be ideal.

A more distinct picture was obtained for the selection scheme: The *egalitarian* variant outperformed both the *elitist* and the *tripartite* versions (also compare Dürr, 2006b; Dürr et al., 2007b). The number of selected individuals has to be chosen such that it balances the improvement intensity and the computational cost. Especially under niching techniques, an ideal number can be broadly expected to be dependent on the multimodality of the problem, since every niche may profit from a local search. Allowing more than 15 to 20 percent of the candidates to be locally altered, however, did not show a significant benefit but only an increase in computation time. On the other hand, a selection rate




Figure 5.12: Comparison of performances of GA and MA on 20-dimensional Ackley function over 25 independent runs, showing maximum, average and minimum best fitness values as a function of the number of evaluations.

of below about 7.5 percent, seemed to undermine the effect of local search. In other words, selection probabilities of 0.075 to 0.15 can be considered ideal for the presented tests.

In all presented tests, the optimization radius was kept constant, with values broadly adapted to the properties of the search space. For example, on Rosenbrock's function a relatively large radius (value of six in all parameters) was employed. There, the difficulty for the optimizer is presented by a non-linear trench in the vicinity of the global minimum. In such a situation, too rigid a restriction may impair the local search by immoderately confining its step length. Similarly, if, like in Ackley's problem, the search space features a multitude of local optima, which obscure the descent direction toward the global optimum, a wider local search radius may be beneficial in that it allows for multiple local optima to be bridged in one local search step. In all other cases a smaller radius on the order of a tenth of the GA-induced coding granularity or less was adopted, mainly for the following two reasons: 1. to limit the computational cost and 2. to restrict the local search to the local attractor in the vicinity of the corresponding candidate obtained by the GA. The latter goal is especially important when using a localized genetic algorithm as in this work. Its achievement, however, is highly idealistic since in general the shape of the search space is unknown and hence a jump to an adjacent niche cannot be prohibited.

Another problem associated with constant optimization radii occurs when the majority of candidates are locally optimal with respect to the radius. Such a situation may be present in late optimization stages. Then, most local search operations will yield optima that coincide with the start solutions. Constructing heuristics to dynamically adapt optimization radii is far from being straightforward, however, and may introduce an unintentional bias on its own.

Finally, the commendable performance of the test using the Nelder-Mead method instead of the SQP is noteworthy; especially in view of the reported problems of this method on high-dimensional search spaces and the simplistic procedure to enforce bound constraints that has been used in this work. The faster convergence rate of the SQP method seems not to countervail the larger number of function evaluations required to numerically determine gradients in the context of memetic search. The situation may be drastically different for (non-linearly) constrained functions, which have not been considered in the presented test suite.

5.6 Local Function Approximation Using Artificial Neural Networks

As argued before, the computation time of fitness evaluations for the problems aimed at in this work can be expected to be high. Consequently, the performance of the optimization approach is chiefly determined by the number of such function evaluations. More drastically, any other routines, including genetic operations or diversity-preserving measures can be considered to be negligible in terms of computation time. The situation is different, however, for the local search step, which involves a number of objective function calls, either to obtain merely zero-order information as in the Nelder-Mead case, or to compute first and second derivatives. To increase its efficiency, the reduction either of the number of function calls or the computation time per call are appropriate measures. A multitude of related ideas applied to evolutionary algorithms have been investigated over the last years. A comprehensive survey of such measures was, for example, given by Jin (2005). There, a broad distinction is drawn between (1) *problem approximations* in which the problem is replaced by a computationally less expensive model, (2) *objective function approximations*, which make use of a simpler surrogate model of the objective function, or (3) *evolutionary approximations*, which targets at reducing the number of actual fitness function evaluations by introducing estimates. Jin (2005) proposes a further subdivision of the last category into techniques he refers to as *fitness inheritance* and as *fitness imitation*. While the former can be broadly described as heuristics under which offspring candidates are assigned fitness approximates derived from their parents, the latter is a method exploiting decision space similarities of solutions in order to perform an educated guess on their fitness. One example is to form clusters of individuals and only rigorously evaluate representatives of these clusters. Other candidates receive the fitness of their equivalence classes. This approach is of course highly sensitive against non-linearities, i.e., epistasis or deception, in the optimization parameters.

In *objective function approximation* techniques, the computation time is reduced by performing evaluations on a model of the fitness function instead of its full-fledged version. For this approach, numerous techniques, well-established in other numerical computation fields, suggest themselves. These techniques (for a discussion, see Jin and Branke, 2005) include response-surface methods or other regression techniques. Alternatively, machine learning approaches such as support vector machines, radial-basis artificial neural networks (ANNs), or, like in this work, multilayer perceptrons (MLPs) are commonly utilized. For the latter to be effective, their capability of serving as universal function approximators is a necessary condition, of course. This has long been a subject of debates, which were often concerned with the question if the superposition theorem by Kolmogorov (1957) is applicable to MLPs. In the late 1980s, several authors independently proved the principal function approximation capabilities of MLPs,²⁰ and later these findings were extended to different neural network variants (for a discussion, see Hassoun, 1995, pp. 46–50). Since then, a number of studies have demonstrated the successful application of ANNs as function approximators in optimization applications (for example, Berke and Hajela, 1992; Carpenter and Barthelemy, 1993), including evolutionary algorithms (for example, Bull, 1999; Pierret and Van Den Braembussche, 1999; Papadrakakis et al., 1999; Jin et al., 2000; Hüsken et al., 2002; Hong et al., 2003). An approach similar to the one here was proposed by Guimarães et al. (2006). There, evolution strategy is combined with SQP. Like in our work, the local optimization radius is restricted to the vicinity of the ES-obtained individuals. Instead of conducting the local search on the original function, a radial-basis approximation, limited also to the proximity of the candidate is generated and used. A user-provided parameter defines the number of generations between local searches, equivalent to our static trigger. In their first version, all individuals of one generation were locally improved on the approximated function. Later, Guimarães et al. (2007) replaced this procedure with a truncation selection scheme, reporting a beneficial convergence behavior when locally optimizing a fraction of a five to ten percent of the population. The result of the local optimization with approximation replaces the original candidate obtained by the ES; no model control is used. An improved convergence rate of the approximated memetic algorithm compared to the non-hybridized ES both on benchmark and on real-world problems was demonstrated by the authors, albeit no comparison with a memetic algorithm without approximation was shown.

²⁰Sometimes referred to as *Cybenko-Hornik-Funahashi theorem* (Cybenko, 1989; Hornik et al., 1989; Funahashi, 1989).

Function approximations using ANNs have been previously applied to lithography simulation. In his approach, which he terms meta-modeling, Capodici (1997) proposes their use as a replacement for time-consuming simulation steps such as PEB (Section 3.4.4). First, a network is trained with a number of pre-computed samples. The so-obtained net is then employed for optimization tasks, also using a genetic algorithm. In contrast to the previous and our fitness function approximation approach, however, candidate solutions are not exploited for a continuous model refinement or control.

Approximation comes at a price. Broadly put, a modification to the objective function or the way the evaluation is performed may have a significant impact on the search behavior of the EA, potentially introducing an unintentional bias. For example, in general, an approximated model cannot be assumed to exhibit the same optimum as the original problem or function. Obviously, the situation aggravates with the required range of validity of the model. As shown in the SQP case (see Section 5.4.1), local models²¹ with a sufficient accuracy and up to a certain confidence can often be found relatively straightforwardly—albeit often requiring extra computation effort. Globally valid models, on the other hand, are rarely available or computable. More importantly, in most cases error estimates or confidence levels cannot be analytically derived at all. To circumvent this problem, it is advisable to set up a regime under which candidates obtained with the approximation model are routinely compared with results from the original formulation. A *model management* has to be established (Jin and Olhofer, 2001). This may be realized, for example, by evaluating the best, average or randomly drawn candidates with the original function to verify the approximation validity and to trigger approximation refinements. As we shall see later, in this work, a slightly different approximation strategy is applied, which could be referred to as *hierarchical fitness model*. There, the fitness evaluation is first performed with a lower accuracy. Only if the resulted merit supersedes a predefined threshold, a finer-grained evaluation is conducted. Especially in the scope of penalty methods, in which infeasible solutions are rejected without regards to their objective value, this approach is suited to drastically reduce the required computation time. Although this approach does not completely evade the problem of an ill-biased search, it does so for the well-performing candidates. Other strategies have been suggested in which the approximation is verified after a constant number of generations, or using dynamic or adaptive intervals.

In this work, we have developed an approximation approach specifically tailored for the local search operation (Popp, 2007). The main reasons for this approach can be summarized as follows:

- (1) The memetic algorithm developed here applies disparate search spaces for the GA (discrete) and the local search (continuous), allowing for a straightforward, non-disruptive extension using an approximation model in the LO step.
- (2) In our approach, local search is followed by a control step, in which the solution is reevaluated in the decision space of the GA. Thus only the exact solution with respect to the GA domain is transferred back into the GA population, and only local modifications outperforming the original solution are accepted. This step implies a model management, rendering approximation model inaccuracies noncritical.
- (3) Principally, the search function needs to be modeled only in the proximity of GA candidate solutions. Therefore the model does not have to account for global properties of the search space.
- (4) An approximation model for the local search step is profitable since depending on the algorithm the local search may require a large number of function evaluations.

5.6.1 Artificial Neural Networks

Like evolutionary algorithms, ANNs are nature-inspired computation models. The basic idea is to leverage a network of relatively simple processing *units*, also called *neurons*, for large-scale tasks by exploiting the high degree of freedom given by the configuration of the interconnections. The resemblance to nervous systems of higher species, often considered the role model of ANNs, is a strictly abstract one (Kecman, 2001, p.13). Albeit there are interdisciplinary efforts to use computer models for the simulation of nervous processes, ANNs are typically merely aimed at computing tasks.

²¹The quadratic replacement problems in SQP.

McCulloch and Pitts (1943) are to be credited for the first concept of neural networks, using binary decision elements and edges (sometimes *synapses*) with fixed weights between these units. In their concept, the units act as threshold gates (see, for example, Rojas, 1996, pp. 27–52). That is, given an input of n bits, the output of the unit is 1, if and only if the sum of the input bits is greater or equal than the threshold θ . Later this concept was extended to incorporating inhibitor edges; for any active inhibition bit, the unit is inhibited, resulting in an output of 0, regardless of the values of the excitatory bits. This extension enables the networks to model any non-monotonic function. When restricting the topology to DAGs, ANNs are termed *feed-forward neural networks*, which because of their earlier discussed principle capability to model any function are suited and hence selected for the purpose of this work. If additionally cycles are permitted, *recurrent neural networks* can be realized, McCulloch-Pitt nets can be used to model any finite automaton. Despite this principal flexibility, a major practical drawback of the former concept is the use of fixed weights of the edges, because in general an adaptation of the network to new functions or automata requires a modification of the thresholds, and more severely, of the topology. In order to circumvent this issue, Rosenblatt (1958) proposed an adaptive-weight model, with variable weights and a fixed topology. Since his original idea was to model a perception/reaction system, this type of ANNs is referred to as *perceptron*. In its simplest form, a perceptron consists of a single layer, with one unit. More formally,

Definition 5.3 (Single-layer perceptron) *A single-layer perceptron is a linear binary classifier, which maps an input vector $\mathbf{x} \in \mathbb{R}^n$ to a boolean value:*

$$y = f_{\theta}(\mathbf{x}; \mathbf{w}) = \begin{cases} 1 & \mathbf{w}^T \mathbf{x} > \theta \\ 0 & \text{otherwise,} \end{cases}$$

where $\theta \in \mathbb{R}$ is an input-independent bias (which can be interpreted as a decision boundary) and $\mathbf{w} \in \mathbb{R}^n$ is a weight parameter vector, where each component corresponds to one input parameter.




Figure 5.13: Single layer perceptron, with inputs $\mathbf{x} : [x_1, \dots, x_n]^T$, weights $\mathbf{w} : [w_1, \dots, w_n]^T$, bias weight $-\theta$, activation function f and output y .

This situation is depicted in Figure 5.13. There, the bias is represented as the constant weight parameter of a constant input of 1, which, as we shall see, facilitates the demonstration of weight-learning algorithms. Because in order for the perceptron to be employed as a classifier for a specific problem, the task remains to find an adequate weight configuration. As indicated in Definition 5.3, single-layer perceptrons are linear classifiers. Thus their classification success rate is one only if the input data is linearly separable; that is, there is a hyperplane that under appropriate weights strictly partitions the input space into the classification targets. Formally,

Definition 5.4 (Linear separability) *Given $X \subset \mathbb{R}^n$ and X_0, X_1 with $X = X_0 \cup X_1$ and $X_0 \cap X_1 = \emptyset$. X_0 and X_1 are called linearly separable : $\iff (\exists \theta \in \mathbb{R} \wedge \exists \mathbf{w} \in \mathbb{R}^n)((\forall \mathbf{x} \in X_0)(\mathbf{w}^T \mathbf{x} \leq \theta) \wedge (\forall \mathbf{x} \in X_1)(\mathbf{w}^T \mathbf{x} > \theta))$.*

The limitation posed by this property is often illustrated by the failure of the perceptron to correctly model the exclusive-or (XOR) gate. As shown in Figure 5.14, because of the non-linear decision boundary, there is no weight configuration allowing for a linear separation.




Figure 5.14: Non-linear-separability of exclusive-or gate results.

The following straightforward learning approach depends on a set of training data (Minsky and Papert, 1969):

$$D := \{(\mathbf{x}, d(\mathbf{x})) \mid \mathbf{x} \in X\},$$

where $X \subset \mathbb{R}^n$ is the input data domain and $d : X \rightarrow \{0, 1\}$ is the desired classification result. Here and in the following, analogous to the bias formulation depicted in Figure 5.13, we assume the input vector to be extended by a constant 1, and the weight vector by $-\theta$, respectively. The error can then be given as:

$$E(\mathbf{x}, \mathbf{w}) := d(\mathbf{x}) - f(\mathbf{x}; \mathbf{w}).$$

In the case of a correct classification, E yields 0, samples falsely classified as 1 or 0 yield 1 or -1 , respectively. An update step is then defined as follows:

$$\hat{\mathbf{w}} := \mathbf{w} + \eta E(\mathbf{x}_l, \mathbf{w}) \mathbf{x}_l, \quad (5.26)$$

where \mathbf{x}_l is an input vector of the training set and $\eta \in [0, 1]$ is a user-provided learning rate.

The update rule (5.26) is applied on the entire training set, and the whole procedure is repeated until a termination criterion, such as a certain error ratio, is met. Novikoff (1962) proved that under linearly separability, this algorithm converges. If the classification sets are not linearly separable, however, the algorithm does, in general, not converge.²²

In order to render ANN learning amenable to gradient descent methods, a differentiable variant of the step function (see Definition 5.3) is commonly used. For that purpose, it is convenient to regard the function f as composition of the *integration function* ($\psi : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$) and the scalar *activation function* (ϕ), i.e.,

$$f_\theta(\mathbf{x}; \mathbf{w}) := (\phi_\theta \circ \psi(\mathbf{w}))(\mathbf{x}).$$

In Definition 5.3, integration is performed as a weighted summation, whereas for the activation a step function was used, which can, for example, be replaced by a (differentiable) sigmoid function,

$$\phi_\theta(x) = \tanh\left(\frac{x - \theta}{2}\right).$$

Then, the correction term of (5.26) is replaced by the following expression:

$$\hat{\mathbf{w}} = \mathbf{w} + \eta(d(\mathbf{x}_l) - f_\theta(\mathbf{x}_l; \mathbf{w}))\phi'_\theta(\psi(\mathbf{x}_l; \mathbf{w}))\mathbf{x}_l, \quad (5.27)$$

²²To circumvent this issue, Gallant (1990) proposed a variant, termed *pocket algorithm* and ensuring convergence by storing the best-so-far observed weight settings.

where ϕ'_θ denotes the first derivative of ϕ_θ . The rationale for this update scheme can be illustrated by assuming the following error due to the current weight configuration \mathbf{w} and with respect to the training instance \mathbf{x}_l :

$$E_D(\mathbf{w}) := \frac{1}{2}(d_l - f_\theta(\mathbf{x}_l; \mathbf{w}))^2.$$

For the gradient with respect to \mathbf{w} we then have²³

$$\nabla E_D = -(d_l - f_\theta(\mathbf{x}_l; \mathbf{w}))\phi'_\theta(\psi(\mathbf{x}_l; \mathbf{w}))\mathbf{x}_l.$$

Applying this error measure to the same training iteration algorithm as before turns it into a gradient descent scheme, often referred to as *delta learning rule*.

The restriction to linear separability of the single layer perceptron can be obviated by employing multilayer networks with multiple processing units (schematically demonstrated in Figure 5.15). For the following considerations, we will formally define the feed-forward neural networks as used in this work. For the sake of a concise treatment, the input parameters are regarded as input nodes, and the outputs net are represented by output nodes.

Definition 5.5 (Feed-forward neural network) *A feed-forward (artificial) neural network is a DAG, with $G = (V_i, V_o, V_n, E)$, where V_i are the input nodes, V_n are the processing units, V_o are the output nodes and E are the directed, weighted edges:*

$$(e = (u, v, w) \in E) \implies (u \in V_i \cup V_n \wedge v \in V_o \cup V_n \wedge w \in \mathbb{R})$$

In order to use the same notation as for the single layer perceptron, a few auxiliary definitions are introduced in addition.

The input to a unit $v \in V_n \cup V_o$, is given by the edges

$$E_v := \{(u', v', w') \in E \mid v' = v\}.$$

The input arity of v is hence $n_v = |E_v|$. Let, like before, $\psi_v : \mathbb{R}^{n_v} \times \mathbb{R}^{n_v} \rightarrow \mathbb{R}$ be the integration function and $\phi_{\theta_v} : \mathbb{R} \rightarrow \mathbb{R}$ the activation function, where θ_v is the threshold corresponding to the node.

Let I be an index set of E_v , with some bijection $s : I \rightarrow E_v$. Then the weight parameter vector is given as $\mathbf{w}_v = (w_i)_{i \in I}$. Further, let $x_u \in \mathbb{R}$ be the output of a node $u \in V_i \cup V_n$ either as the result of the activation function ϕ or, if $u \in V_i$, an input parameter. Then, we can define the input of v as $\mathbf{x}_v = (x_i)_{i \in I}$.

In addition to misplaced folklore, which arrived at the conclusion that the results for the single-layered perceptron are valid also for multilayer variants, a major issue with multilayer variants was the lack of adequate learning algorithms. The popularity of ANNs drastically increased in the mid-1980s with the development of such learning algorithms. Today, a large number of different approaches are available, which can be distinguished according to the following categories:

- *Supervised learning* involves training data that is presented to the neural network. As shown for the delta rule, which also belongs to this category, the training data consists of input parameters and the corresponding target response of the neural net. Both the selection of the training set and its presentation to the network are normally managed by an external agent. Weights are adapted such that a satisfactory agreement between target and actual response is achieved.
- In the case of *unsupervised learning*, no teaching agent is involved. Instead, the network is expected to autonomously define rules based on the data it is confronted with. A common application is to reveal (a priori unknown) structures of data, for example, by clustering.
- *Reinforcement learning* is commonly applied in tasks for which an ad hoc response to a, potentially dynamic, environment is required. There, learning is performed implicitly, by adapting the weights based on the consequence of the current response, often through a reward mechanism. In addition, the exploration of the search space is autonomously (e.g., heuristically) performed by an agent. In a way, evolutionary algorithms can be said to utilize the reinforcement learning paradigm.

²³For more details, see Rojas (1996, chap. 7)




Figure 5.15: Multilayer perceptron with an input layer, two single hidden layer and an output layer. Corresponding nodes are from the sets V_i , V_n and V_o , respectively. The weight edges (connections) are given by a triple (u, v, w) , where u and v are the connected nodes and w is the weight factor of the connection. The inputs to each node are denoted as a vector \mathbf{x} , the corresponding weights as \mathbf{w} . Each processing unit (nodes in the hidden layer) consists of an integration function ψ and an (typically sigmoid) activation function ϕ .

In this work, for the task of function approximation, we have restricted our focus to supervised learning strategies, for which a common approach is the backpropagation algorithm (Rumelhart et al., 1986; Werbos, 1990). It can be considered a recursive extension of the gradient descent approach for multi-layered feed-forward networks. Since only at the output layer, the target pattern can be provided—hidden layer output can be inferred, but no desired state can be defined—an error can be determined only there. The basic idea is to propagate this error backward through the network, applying the same weights as in the forward direction. The principal flow for a net with one hidden layer is shown in Algorithm 5.3. Similarly as before, the training set is defined as,

$$D := \{(\mathbf{x}, \mathbf{d}(\mathbf{x})) \mid \mathbf{x} \in X\},$$

where $\mathbf{d}(\mathbf{x}) = (d_v(\mathbf{x}))_{v \in V_o}$ are the target responses due to \mathbf{x} .

Algorithm 5.3: Backpropagation Learning for ANN with one hidden layer

```

/* randomly initialize all weights */
1 for (u, v, w) ∈ E do
2   | w ← random initialization
3 while termination criterion not met do
4   | /* for all training samples */
5   | for (x, d(x)) ∈ D do
6   |   | /* feed-forward step */
7   |   | foreach v ∈ Vo do
8   |   |   | ov(x) ← output due to x
9   |   |   | δv ← ov(x) − dv(x)
10  |   | /* backward step */
11  |   | /* compute hidden layer to output layer errors */
12  |   | foreach u ∈ Vn do
13  |   |   | δu ← 0
14  |   |   | foreach (u, v, w) ∈ E do
15  |   |   |   | δu ← δu + wδv
16  |   | /* compute input layer to hidden layer errors */
17  |   | foreach u' ∈ Vi do
18  |   |   | δu' ← 0
19  |   |   | foreach (u', v', w') ∈ E do
20  |   |   |   | δu' ← δu' + w'δv'
21  |   | update weights

```

For a feed-forward network with a fixed topology, in which each node of a layer is connected with all nodes of the subsequent layer, an alternative concise notation can be adopted (cf. Rojas, 1996, pp.170–171). Let $\mathbf{W}^{(l)}$ be the weight matrix from layer l to layer $l + 1$, with

$$\mathbf{W}^{(l)} := (w)_{i,j}^{(l)},$$

where i denotes a unit in layer l that connects to a unit j in $l + 1$. Thus given m units in l and n in $l + 1$, we obtain an $m \times n$ matrix. Given the derivatives of the activation functions at layer l , which can be computed and stored during the feed-forward step, as the diagonal matrix:

$$\mathbf{D}^{(l)} := \text{diag}((d)_i^{(l)}),$$

for example, with

$$d_i^{(l)} = \phi_i^{(l)}(1 - \phi_i^{(l)}),$$

in the case of a sigmoid function. Then, if we define ϕ as the output of the neural network due to the current weights and the presented input, and, like before, \mathbf{d} as the target pattern, we have that

$$\delta^{(\ell)} = \mathbf{D}^{(\ell)}(\phi - \mathbf{d}),$$

and we obtain the following recursion formula:

$$\delta^{(l)} = \mathbf{D}^{(l)} \mathbf{W}^{(l+1)} \delta^{(l+1)} \quad l = 1, \dots, \ell - 1.$$

Given an additive integration function as before, we can infer the weight correction matrix for a layer as follows:

$$\Delta \mathbf{W}^{(l)} = \eta \delta^{(l)} \left(\mathbf{x}^{(l-1)} \right)^T.$$

The above scheme can be readily used in Algorithm 5.3. There are two different modes in which the backpropagation algorithm can be executed:

- (1) As an *incremental learning* routine, in which one or a few training instances are presented to the network, leading to a instant weight correction, transforming the approach into a stochastic gradient descent procedure,
- (2) or as a *batch learning* routine, in which the entire training set is used for one iteration, rendering this approach a classical gradient descent routine.

Incremental learning offers a number of advantages (LeCun et al., 1998), especially with respect to the performance of learning since training samples are typically only considered in case of large errors and can be omitted if the observed output already shows a sufficient agreement with the presented target. Furthermore, stochastic gradient descent methods are generally less prone to prematurely converge to stationary points that are due to the form of the activation function or noisy training patterns. The non-determinism so met with, however, is also a source of error, potentially degrading its accuracy and, more importantly, its precision.

To stabilize convergence, especially in search space regions of ill-directed gradients, it is common practice to average the descent direction over a number of iterations (epochs). This technique sometimes referred to as *gradient descent with momentum* since it introduces an additional parameter, called momentum. Momentum defines the fraction by which a previous candidate contributes to the current gradient determination (Rumelhart et al., 1986). Alternatively, traditional line-search has been proposed to rigorously adapt the step width (Jones et al., 1990).

Another option to improve the convergence behavior of backpropagation is the use of adaptive learning rates. In that spirit, the resilient backpropagation (Rprop) (Riedmiller and Braun, 1992) makes use of step lengths, which are determined individually for each weight. Specifically, the learning rate is decreased if the sign of the partial error derivative with respect to the weight has changed in the current epoch. It is increased, in the case of a non-changed sign, and unaltered in case of a derivative of zero, which as we shall see is for algorithmic convenience. First we define,

$$\Delta \partial E_{i,j}^{(k)} := \left(\frac{\partial E}{\partial w_{i,j}} \right)^{(k-1)} \left(\frac{\partial E}{\partial w_{i,j}} \right)^{(k)}$$

as the according derivative change from $k-1$ to k . We thus have for the learning rate adaptation that

$$\eta_{i,j}^{(k)} := \begin{cases} \eta^+ & \Delta \partial E_{i,j}^{(k)} > 0 \\ \eta^- & \Delta \partial E_{i,j}^{(k)} < 0 \\ 1 & \text{otherwise} \end{cases} \eta_{i,j}^{(k)}, \quad (5.28)$$

where $\eta^- \in [0, 1)$ and $\eta^+ > 1$. Typically, the new learning rate is also bound to some $[\eta_{\min}, \eta_{\max}]$. In addition to a learning rate adaptation, the original Rprop version uses a weight backtracking mechanism. A number of different variants have been proposed, we will follow the classification of Igel and Hüsken

(2000) in terming (1) The original weight-tracking version Rprop⁺, (2) the variant without weight-tracking Rprop⁻, (3) the improved weight-tracking version iRprop⁺, (4) and the improved, non-tracking variant iRprop⁻.

The step length part of the Rprop algorithm is depicted in Algorithm 5.4. There, both variants Rprop⁺ and Rprop⁻ are illustrated (backtracking is controlled through the flag *weight-tracking*). The improved variant, devised by Igel and Hüsken (2000), incorporates a minor change, which according to benchmark results provided by the authors, leads to a significant performance improvement. It applies a backtracking step only if, in addition to a change of the error derivative direction, the global error increases due to a given weight change. Broadly put, for a situation in which the weight configuration is close to an optimum, for which the total error can be regarded as a criterion, an oscillation of the descent direction may still lead to a faster convergence compared to a step length optimization under an enforced descent direction. Consequently, the authors propose to conduct the retraction step (Algorithm 5.4; 1.7) only if additionally

$$E^{(k)} > E^{(k+1)}.$$

Algorithm 5.4: Resilient backpropagation algorithm (Rprop)

```

1 foreach  $w_{i,j}$  do
2   if  $\Delta\partial E_{i,j} \neq 0$  then                                     /* no backtrack step */
3      $\eta_{i,j} \leftarrow$  update learning rate                       /* as in (5.28) */
4     if  $\Delta\partial E_{i,j} \geq 0$  then                                  /* same derivative direction as in previous epoch or
       backtracking */
5        $w_{i,j} \leftarrow w_{i,j} - \begin{Bmatrix} 1 & \partial E/\partial w_{i,j} > 0 \\ -1 & \partial E/\partial w_{i,j} < 0 \end{Bmatrix} \eta_{i,j}$ 
6     else if weight-tracking then                             /* derivative direction changed */
7        $w_{i,j} \leftarrow w_{i,j}^{(k-2)}$                              /* retract weight change */
8        $\Delta\partial E_{i,j} \leftarrow 0$ 

```

A closer local approximation of the error function with respect to some weight settings can be achieved by incorporating second-order information. Along that line, Fahlman (1988) has proposed a backpropagation variant that he termed *QuickProp* and which uses quadratically approximated error responses. Instead, however, of actually computing second derivatives, a secant approximation is employed. The routine can hence be considered a pseudo-Newton method. For the weight correction for an iterate k , we have

$$(\Delta w_{i,j})^{(k)} := (\Delta w_{i,j})^{(k-1)} \frac{\frac{\partial E}{\partial w_{i,j}^{(k)}}}{\frac{\partial E}{\partial w_{i,j}^{(k-1)}} - \frac{\partial E}{\partial w_{i,j}^{(k)}}}.$$

As in the traditional backpropagation approach, first-order derivatives can be conveniently computed during the feed-forward step.

Numerous other backpropagation variants have been proposed, including second-order methods, for example, using BFGS and low-memory BFGS Hessian approximations, Levenberg-Marquard methods (Battiti, 1992), or evolutionary algorithms (Yao, 1993). The performance of these methods strongly depends on the ANN application (where the network topology is an important factor). Because of the proven track record of successful applications and their wide availability within ANN frameworks, we have restricted our investigations to classical backpropagation, iRprop and QuickProp.

5.6.2 Incorporation of Function Approximation for Local Search

In order to conduct the local search on a function approximated by a neural network, the respective facilities have to be integrated into the memetic algorithm. Specifically, (1) the training data set has to be inferred, stored and provided for the training phase of the ANN; (2) the local search has to be adapted to operate on the approximated instead of the exact function; and (3) a model control, verifying the quality of the local-search-obtained candidate on the exact function has to be established.

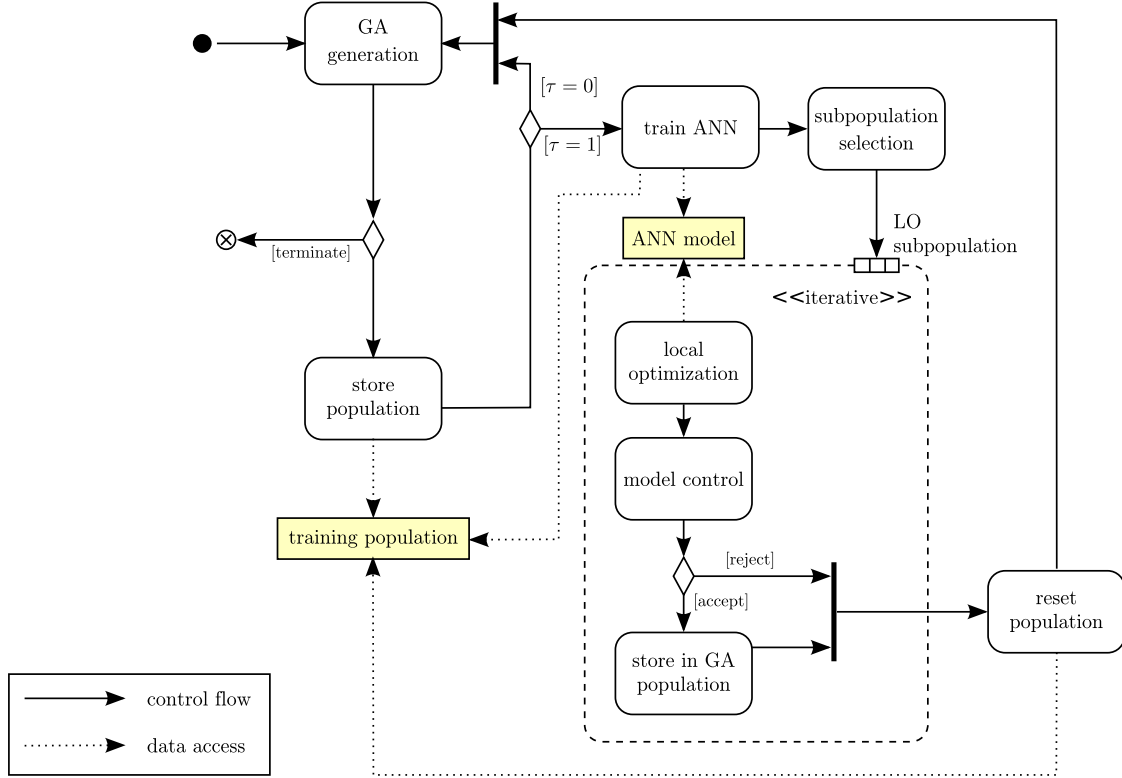


Figure 5.16: Schematic of flow of memetic algorithm with function approximation.

The according flow of the memetic algorithm with fitness approximation is depicted in Figure 5.16 (modified after Popp, 2007). After each GA generation, the current population $P^{(k)}$ is added to the training population (Popp, 2007):

$$P_T := P_T \cup P^{(k)}.$$

Additionally, like in the non-approximated memetic algorithm (MA), a trigger test is performed, determining whether a local optimization is to be conducted or not. Here, a *static trigger*, with an interval of 15 generations is used. In case of a negative trigger result ($\tau(k) = 0$), another GA generation is performed; otherwise the individuals for the local search step are selected. In accordance with the previously discussed tests, a *random selection*, with a probability of 0.1 is performed. In other words, about ten percent of the current population are subjected to local improvements. Let $P_l \subset P^{(k)}$ be the selected sub-population. The stored population P_T is then used to train the feed-forward ANN. The training—incremental or batch—is terminated after a sufficiently small error was achieved or a maximum number of epochs has been reached. The ANN model of the function is used as the objective function of the subsequent local optimization. We have restricted our analysis to the Nelder-Mead down-hill simplex algorithm, which we have found to yield acceptable local improvements. All candidates of the local selected sub-population P_l are then locally optimized. As before, the continuous representation of the parameters current decision space parameters are used as the start value for the local optimization. No further transformation as in the case of the MA without approximation is required. Since the ANN

yields output values in the range of $[-1, 1]$, the fitness has to be rescaled according to the range of the objective space (ibid.). Rescaling can be performed as a step in the fitness calculation, i.e.,

$$\mu : \phi \mapsto \frac{1}{2}((\phi + 1)y_u + (1 - \phi)y_l)$$

where ϕ denotes the approximated objective function and $[y_u, y_l]$ is the range of values to be obtained by μ . Equivalently, the scaling can be applied to the final solution of the local optimizer. Since the actual objective value is discarded in the following steps, however, also scaling can be omitted. Instead, the locally optimized point undergoes a model control step. For that purpose, its parameters are transformed back into the MA decision space representation, which like before is taken to be the binary coding of the GA. The individual is then evaluated using the same fitness function employed in the GA. As before, in the case of unconstrained local optimizers such as the NM method, a clipping of box-constraint violating parameters is performed. The control evaluation can be expected to be the computationally most expensive step. It serves two purposes: first, it determines if the local, approximate fitness improvement prevails (a) on the exact fitness model and (b) using the nominal decision space representation. Particularly, if under exact fitness evaluation the locally modified individual does not outperform the original, GA-obtained solution, it is rejected. Only in the case of an improved fitness, the solution replaces the original candidate. Second, as in the non-approximate MA, the individual is assigned the exact fitness value, which serves as the criterion for the resumed GA and its operations. After all individuals of the local subpopulation have been optimized and controlled, and all accepted candidates have been transferred back into the GA population, the training population is deleted. This is done so in order to adapt the approximation model to the current state of the GA population and the according figure of the search space (ibid.). It could be broadly described as an ad hoc, snapshot-based model generation technique and is motivated by the conjectured properties of the so-obtained model: It accentuates the same search space regions as the GA does at the time of model generation. In other words, the model properties loosely follow the GA stages. In an early phase, in which the population is sampled rather arbitrarily, it can be expected to show a coarse global approximation. Later stages, can be assumed to lead to a model with a finer approximation of certain search space regions, because there exploitation plays a larger role. Similarly, if niching techniques are made use of, the ANN model should show an increasing adaptation to the emerging niches.

Similarly as for the function values, a scaling for the optimization parameters has to be devised, which maps the input to the $[-1, 1]$ interval. This could be done in a straightforward way by exploiting the box constraints of the GA. As Popp (2007, chap. 3) points out, however,²⁴ this could lead to a situation in which all parameter values are either negative or positive, which would drastically impair the learning process. He thus devised the following scheme: Given an optimization parameter x and the mean value \hat{x} due to the current generation. Further let x_l, x_u denote the current minimum and maximum values, respectively. The maximum absolute deviation from the mean is obtained by,

$$x_d = \max(|\hat{x} - x_l|, |x_u - \hat{x}|),$$

such that $[\hat{x} - x_d, \hat{x} + x_d]$ constitutes a symmetric range. And hence

$$[\hat{x} - x_d, \hat{x} + x_d] \rightarrow [-1, 1], \quad (5.29)$$

yields the desired property of a mean of zero of scaled parameters.

As argued before, ANNs with at least one hidden layer are suited for universal function approximations. The total number of processing units can of course be expected to somehow reflect the “complexity” of the search space, for example, such as multi-modalities. On the other hand, for a efficient learning phase, the number of connections and consequently the number of nodes is to be kept as small as possible. For the purpose of this work, the number of input and output nodes are predetermined: there are as many input nodes as there are optimization parameters, and, since we are concerned with a single-objective MA, a single output node is employed.

²⁴following LeCun et al. (1998)

The performance of the learning algorithm strongly depends on the network topology, on the search space and on the training samples. Therefore an analytic selection of the algorithm and its parameters, such as the learning factor, is generally not possible. We have hence conducted several numerical experiments, from which we try to estimate rules for ideal settings. Of course, this procedure is rather imprecise in that it will not avert problem-specific adaptations. However, since precision is not a top priority for the local modification operator, it may result in good enough a heuristic for the real-world applications aimed at.

5.6.3 Benchmark Problems

In order to infer suited ANN configurations, including topology, and the learning algorithm and its parameters, a number of benchmark tests were conducted (Popp, 2007, chap. 4). The goal of the training phase was to achieve a mean-squared-error (MSE) of 0.001 or smaller, which in numerical pre-studies was identified as an adequate compromise between learning efficiency and network quality (ibid.). The first benchmark was conducted with the Rosenbrock function (see Appendix F), with 15 and 30-dimensional inputs.²⁵ The GA population sizes were set to 150 for the 15-dimensional and to 300 for the 30-dimensional case; maximum generations were 500 and 1500, respectively. The network topology was determined empirically by evaluating the learning success, the achievement of the designated MSE for different configurations. Both one and two-hidden-layered nets were considered. As the training algorithm, stochastic backpropagation with momentum (α) of 0.9 and a learning rate (γ) of 0.2 was applied. The MA was set up to use a constant trigger, executing local search every 15 generations, for which the net was trained with the preceding generations. Thus in the case of 500 generations, 33 ANN approximations are generated. Egalitarian local search selection with a 0.1 selection probability was employed. A canonical, generational, binary-coded GA was used for this benchmark study, with point mutation ($p_m : 0.05$), two-point crossover ($p_c : 1.0$), roulette wheel selection, and without elitism.

Table 5.1: Comparison of incremental learning results for different network topologies on the 15-dimensional Rosenbrock test function. The first row shows the topology; for two-hidden-layered nets, nodes per layer are separated by a dash. The success rate, that is, the number of cases in which the net achieved an error equal to or less than the target within the maximum number of 1000 epochs is presented in the second row. The last row contains the average number of epochs required to achieve the targeted MSE in the success case.

Hidden nodes	16	18	20	7-7	8-8	9-9
Success rate	0.93	0.95	0.95	0.67	0.81	0.95
Avg. number of epochs	452	479	443	556	578	429

As shown in Table 5.1 for the 15-dimensional case, topologies with a single hidden layer exhibit a low sensitivity against the exact number of nodes both with respect to the training success and the number of epochs required to achieve the error target. For the two-layered case, the minimum node number for an acceptable learning success rate seems to be 9-9 (Popp, 2007, chap. 4). Leaner configurations drastically degrade the success probability while at the same time increasing the average epochs. In the case of a single-hidden-layered net, the exact number of nodes have a less pronounced impact on the learning success and the learning speed can be observed (ibid.). The approximate rule, proposed by Berke and Hajela (1992) and Carpenter and Barthelemy (1993), which suggests a number of hidden-layer nodes on the order between the average and the sum of input and output nodes, appears to also be applicable here.

In a next step, we have evaluated the impact of the algorithm on the problem-specific performance of learning (see Popp (2007)). Taking the same Rosenbrock function example, we have conducted an

²⁵Here, the function is multiplied with -1 to obtain a maximization problem.

Table 5.2: Comparison of learning success for different algorithms and parameters on Rosenbrock's function for a single hidden layer with 20 nodes and for two hidden layers with nine nodes each (9-9). Incremental backpropagation is conducted with different learning rate (γ) and momentum (α) configurations; for QuickProp, the learning rate is varied, while iRprop is a parameter-less algorithm. Each result consists of two values, the average success rate and the average number of epochs in case of a successful training, in which the designated MSE was achieved. The best results are typeset in bold. The worst results are shown in italics.

Hidden Layers	20	9-9	20	9-9	20	9-9
Incremental Backpropagation						
$\alpha \backslash \gamma$	0.2		0.5		0.7	
0.0	0.97	0.91	0.95	0.92	0.95	0.95
	<i>482</i>	463	434	434	443	439
0.5	0.96	<i>0.87</i>	0.95	0.93	0.95	0.93
	438	431	463	441	462	437
0.9	0.95	0.95	0.93	0.9	0.95	0.92
	443	429	432	478	424	479
QuickProp						
γ	0.2		0.5		0.7	
	0.95	<i>0.88</i>	0.92	0.95	0.96	0.89
	442	442	478	416	460	<i>483</i>
iRprop						
	0.95	0.96				
	405	462				

experiment in which we compared incremental backpropagation, QuickProp, iRprop and the respective parameters. The results are summarized in Table 5.2. Again, the 20-node, single-hidden-layer network proved relatively insensitive against the learning algorithm and parameter selection. The best result in terms of required learning epochs was achieved with the parameter-less iRprop approach, where at the same time the success rate was acceptable. Critical success rates below 0.9 can be observed only for the two-hidden-layered net.




Figure 5.17: Comparison of performances of approximated and non-approximated memetic algorithm on 15-dimensional Rosenbrock function. The solid line shows the result obtained with a one-hidden-layered 20-node ANN, the dashed line demonstrates the convergence for a two-layered ANN with nine nodes in each layer, and the dotted line denotes the MA without approximation. Both fitness values and number of evaluations are shown in logarithmic scale.

Figure 5.17 shows a comparison of the convergence rate obtained by the two approximated (20 and 9-9 hidden-layered) and the non-approximated MA, all other settings identical. Specifically, both algorithms employ the same Nelder-Mead technique, limited to a maximum of 50 iterations. The initial simplex was set to a 0.001 expansion in all directions. The horizontal axis in the graph denotes actual function evaluations. That is, calls of the approximated function are not shown. All results constitute an average over five independent runs. Neither the single-layer nor the two-layer approximation models reach the same sub-optimal solution as the MA.²⁶ The flat plateau in the vicinity of the minimizer of Rosenbrock’s function render it difficult for both the GA and the Nelder-Mead algorithm. The approximation model aggravates the situation in that even small local deviations from the exact objective function may lead to a different local search behavior. To illustrate, we have conducted an experiment using Rosenbrock’s function in two dimensions (Popp, 2007, chap. 4). There, the domain was restricted to $[-10, 10] \times [-10, 10]$.

Figure 5.18 shows the graph of the exact along with the approximated function, which was obtained from the iRprop-trained ANN after the first 15 GA generations, for which candidates in a domain of about $[-8, 8] \times [-10, 10]$ occurred. The corresponding, non-scaled approximated function is given by $f : [-1, 1]^2 \rightarrow [-1, 1]$. In addition to the apparent errors at the bounds of x_1 , the approximate shows a small, yet effective deviation within the plateau, and hence to an optimum different from the nominal one at $(1, 1)$.

A similar situation occurs in the 30-dimensional Rosenbrock case. For example, conducting an experiment with 30 input nodes and a single-hidden-layer with 25 nodes, again using the iRprop training algorithm, and five independent runs, shows a close agreement of the resulted sub-optimum with the

²⁶Although, also in the approximation case, the found parameters $\hat{\mathbf{x}}$ are close to the minimizer \mathbf{x}^* , with $x_i^* = 1$ ($i = 1, \dots, 15$), i.e., $|\hat{\mathbf{x}} - \mathbf{x}^*| < 2$.




Figure 5.18: Comparison of exact graph of two-dimensional graph of Rosenbrock’s function and the approximate obtained through ANN training after 15 MA generations. Approximate function is presented in the ANN domain, i.e., $f : [-1, 1]^2 \rightarrow [-1, 1]$.

one obtained through the non-approximated MA (see Figure 5.19).²⁷ The high dimensionality of the problem, however, presents a pronounced problem to the NM method, which is reflected by the large average number of function evaluations (912) required per local improvement step.

In order to verify the performance of the proposed approach on multimodal search spaces, another test was performed on Schwefel’s function (Appendix F), with a 20-dimensional decision space (Popp, 2007, chap. 4). The population size was 150. The maximum generation number was 1500. In this test, incremental backpropagation using a momentum of 0.9 and a learning factor of 0.2 was applied, where the same MSE of 0.001 was aimed at (and on average reached after 196 epochs). As depicted in Figure 5.20, which shows the convergence graph for five independent runs, the convergence of the MA with approximation is more rapid—with respect to the required function evaluations—than the non-approximated one. However, the search performed on the approximate model leads to an oscillation around a fitness of about -0.9, which is above the average value of -0.013 found with the exact model. The corresponding differences in the decision space are small. For example, the (Euclidean) distance between the global minimizer and the solution found by the exact MA is about 0.27, while global minimizer and the average solution of the approximated MA are separated by a distance of less than three. In other words, both MAs have found solutions, whose attractor is the global minimum. The imprecisions can be attributed to model and sampling inaccuracies and their interaction.

In conclusion, the benchmark experimentations have demonstrated the performance improvement potential of the proposed approximation scheme in conjunction with the MA. Even for the highly multi-modal search space of Schwefel’s function, no significant degrade of the accuracy of the found solutions occurred. Similarly, in the Rosenbrock function case, which because of its quasi-stationary region can be considered “hard” for both the evolutionary algorithm and the NM method, the inclusion of an approximation, did not lead to a deteriorated convergence. It can be speculated (but has not been verified here) that an intensified search or a meta-MA approach, as explained earlier, may exhibit further improvement potentials. Of course, the actual performance strongly depends on numerous other actuating factors set by the problem and the local and global search configuration. For example, niching

²⁷In contrast to the 15-dimensional case, the population size was increased to 500, and the maximum number of generations was 1500.




Figure 5.19: Comparison of performances of approximated (solid line) and non-approximated (dotted line) memetic algorithm on 30-dimensional Rosenbrock function. For the ANN, a single hidden layer with 25 nodes was used. Both axis are in logarithmic scale.




Figure 5.20: Comparison of performances of approximated (solid line) and non-approximated (dotted line) memetic algorithm on 20-dimensional Schwefel function. (a) Entire convergence graph in double logarithmic scale. (b) Close-up of convergence graph around 10 000 evaluations. The non-approximated MA shows a significantly longer retention time per fitness value. This is because of the high number of function evaluations required for the local search on the exact function.

approaches, not considered in the former tests, can be expected to concentrate the approximation model to locally sub-optimal regions, posing both a better approximation in the respective regions but also the risk of an overfitting of the ANN and hence a loss of fidelity.

6 Simulation and Optimization Framework

The feasibility of the optimization approach aimed at in this work depends on a number of aspects, including an efficient lithographic simulation environment, the implementation of the employed optimizers and mechanisms to exploit the inherent parallelism of the evolutionary algorithms. This chapter presents a brief overview of the respective tools.

6.1 Lithography Simulation Environment

As discussed in Chapter 3, the simulation of the lithographic process requires a wide range of numerical models, at least a subset of which can also be expected to be involved when performing process optimizations. For the examples presented in the next two chapters, the following solvers are essential:

- A rigorous EMF solver, applicable both to the mask and the wafer domain.
- An efficient vector imaging model, which supports incidence-dependent spectra and captures thin-film effects.

In addition, the optimization routine needs to be able to efficiently execute the employed solvers and to access the resulted data. The simulation framework *Dr.LiTHO* (Erdmann and Henke, 1999; Fühner et al., 2007b) used in this work meets these requirements by providing:¹

- An FDTD solver (see Section 3.1.1) and the waveguide method, including a domain-decomposition method (see Section 3.1.2).
- An Abbe imaging routine (Section 3.2.5) based on both fixed source and pupil grids. Furthermore, a more recent variant (Evanschitzky et al., 2009) allows for an arbitrary source point specification and continuous-domain projector pupil definition. Both implementations employ vector imaging models, support incidence-angle-dependent spectra and are able to account for thin-film effects.
- An interface to the purely interpreted programming language Python that enables full access to most of the internal data.
- Photoresist models as discussed in Section 3.4.

6.2 Single-objective Genetic Algorithm

The single-objective genetic algorithm used in this work is implemented in C and comprises a Python interface. This architecture permits a direct integration with the previously described lithography simulation environment. For example, fitness functions can be tailored for a direct interaction with the simulation environment. Furthermore, the object-oriented design of the GA, albeit not using an object-oriented language for its core, allows for a straightforward extension, either by adding additional components to its C core or, for less involved modifications, by expanding the Python interface. To highlight its inherent niching capabilities—using restricted tournament selection (RTS) (see Section 4.9.2)—this implementation has the recursive acronym “GiNGA is a niching genetic algorithm (GiNGA).” GiNGA is aimed mainly at being a binary-coded GA, supporting standard binary and Gray code. Real-coded chromosomes can be used as well but are not considered in this work. For the optimization problems examined here, a number of different domains are employed, including integers, real numbers or bit patterns, for which also different binary coding schemes are devised. The standard mapping for real numbers can be briefly described as follows: Given a parameter $x \in \mathbb{R}$ and a box constraint:

$$l \leq x \leq u,$$

¹The Development and Research Lithography (Dr.LiTHO) simulation environment contains a large number of additional modules and models not discussed here.

$l, u \in \mathbb{R}$. Additionally, let d denote the number of desired decimal places, by which the target granularity or precision of x is determined. The required bit-string length for the parameter x is then given by

$$\ell := \lceil \log_2 ((u - l)10^d + 1) \rceil,$$

where $\log_2(\cdot)$ denotes the logarithm of base two and

$$\lceil u \rceil := \min\{z \in \mathbb{Z} \mid z \geq u\} \quad (u \in \mathbb{R}).$$

Consequently, the (truncated) binary representation of x is found by mapping the number to the range $[0, 1]$ and by then scaling it according to the previously obtained bit-string length:

$$a_{10} := \frac{x - l}{u - l}(2^\ell - 1).$$

Finally, this number is converted into the binary or Gray representation. Vice versa, to decode a single parameter after its position in the bit-string is located and its decimal representation is obtained, denoted again by a_{10} , we have that

$$x := \frac{a_{10}}{2^\ell - 1}(u - l) + l.$$

GiNGA provides the following operators:

- (1) *one-point*, *two-point* and *uniform* crossover (see Section 4.6),
 - (2) *point mutation* (see Section 4.7) and
 - (3) *roulette wheel*, *rank-based* and *tournament* selection (see Section 4.5),
- employing the usual control parameters (see Section 4.8). Parameter adaptation have not been considered in this work. GiNGA is a generational, elitist GA, where the number of elite solutions can be user-configured. Where not otherwise noted, the following settings have been used:

- Gray code,
- restricted tournament selection,
- two-point crossover (p_c : 0.6),
- point mutation (p_c : 0.005) and
- no explicit elitism under RTS.

GiNGA has been used in a number of applications not related to this work (Tollkühn et al., 2003; Fühner and Jung, 2004; Ortiz et al., 2004).

6.2.1 Caching of Fitness Values

In most optimization applications, the evaluation of a solution's merit is the most expensive step in terms of computation time. In contrast, genetic operations, such as crossover or mutation, require only a very small fraction of the overall runtime. On the other hand, some GA operations require a repeated evaluation of individuals for which a fitness has been previously determined. Although some of these operations can be implemented such that the GA is aware of valid merit values, in many cases this is either not possible or would involve a high level of administrative routines. For the following examples no further bookkeeping is required:

Steady State GAs In the case of steady state GAs, only the worst individual of the current population is replaced by a newly created offspring. The rest of the population is left unaltered. For these remaining solutions no further fitness calculations are required.

Elitist Selection As with steady state GAs, for elitist selection the individual's chromosome does not change at all. Thus, no new fitness evaluation is performed.

Recombination When performing recombination, only a certain percentage (specified by the crossover probability $p_{\text{crossover}}$) of the selected individuals are actually mixed. The other solutions are simply copied. The fitness for the copied individuals remains unchanged.

A moderate administrative effort has to be made for RTS. With RTS, individuals are randomly chosen and recombined. Obtained offspring individuals then replace their closest neighbors in the former population if and only if their fitness is higher than the neighbors' fitness. For this selection operation, fitness values might have to be re-calculated several times. However, it is relatively simple to add a mechanism that identifies if a fitness value is still valid or needs to be recalculated because of an interfering genetic operation (e.g., crossover).

In addition to the former cases, identical individuals may also be generated through other mechanisms. Especially in a later state of a GA run, with high selection pressure the probability of equivalent individuals as a result of genetic operations increases. In order to also exploit pre-evaluated fitness values for these cases, it is necessary to store merit results independent from the individuals. For that, a caching mechanism was developed: Before evaluating an individual's pay-off by calling the objective function, it is first tested whether a merit corresponding to the solutions chromosome is already stored. If so, this value is assigned to the new solution. If it is not present in the cache, the fitness function is called, the yielding value is assigned to the individuals but is also stored in the cache. Algorithm 6.1 demonstrates the flow.

Algorithm 6.1: Fitness Cache

```

1 if IsInCache(A) then
2    $\mu(A) \leftarrow \text{GetFitnessFromCache}(A)$ 
3 else
4    $\mu(A) \leftarrow \text{Evaluate}(A)$ 
5   StoreInCache(A,  $\mu(A)$ )

```

For binary strings with fixed lengths (which are common for GAs), the cache can be organized very efficiently by regarding the chromosome as a binary tree. Fitness values are stored as leaves of this tree. Figure 6.1 gives an example of this tree representation of a chromosome with length three, yielding eight (2^n , where $n = 3$) different chromosome combinations. Instead of having to linearly search through a maximum of all eight stored chromosomes, this representation limits the maximum search depth to three ($n = \log_2(2^n)$). Each bit can be regarded as a node determining whether to follow the left or right branch. In the final node or leaf, the fitness value is stored. An incomplete path indicates that no fitness value corresponding to the chromosome is found. In this case, the fitness value has to be actually evaluated.




Figure 6.1: Binary tree representation of chromosomes.

Objective values are stored in the same way. Non-existing branches are created such that a complete path, reflecting the chromosome's bit, down to the leaf level is generated. Even for very long bit-strings and a large number of chromosomes, this strategy is efficient and helps to save numerous time-consuming calculations.

Example 6.1 Given a chromosome length of 50: The number of possible combinations is hence $2^{50} = 1125899906842624$. However, the maximum depth of the binary tree is 50. Thus, in order to retrieve a fitness value from the cache, a maximum of 50 branch-tests has to be performed. Also, in case of a cache miss, that is, the corresponding fitness was not saved, the time loss corresponds to 50 branch-tests.

6.3 Multi-objective Genetic Algorithm Implementation

The multi-objective GA adopted in the scope of this work is termed *Python multi-objective evolutionary algorithm (Pythmea)* and was developed by Seifert (2006a). It implements both the NSGA-II (see Section 4.11.3) and the SPEA2 (see Section 4.11.4) algorithms. Its modular concept, however, also allows for the recombination of different aspects of different algorithms, for example, combining NSGA-II with an SPEA2 archive. Like GiNGA, its default representation is binary, employing the same standard coding scheme for integer and real numbers. Pythmea uses the same GA operators and defaults as GiNGA. Additionally, it provides functions to analyze optimization results and to assess the optimization performance (see Section 4.11.5).

6.4 Problem Definition Interface

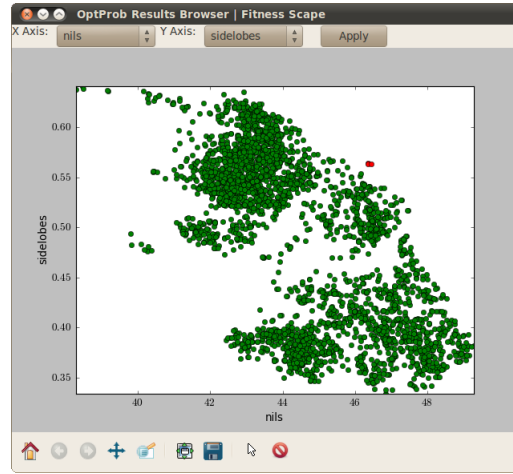
For an optimization approach to be efficient, the problem formulation has to be closely tailored to it. Depending on the assumptions made and on the numerical models employed, an exchange of the optimization approach is often far from being straightforward. This situation is aggravated by the fact that pre-existing optimization routines often greatly differ in terms of their interface. That is, even if a problem could be principally effortlessly reformulated to render it amenable for a different optimizer, a number of practical, implementation-related problems and incompatibilities may lead to significant complications. In this work, a single problem is often solved with multiple optimization approaches, for example, to study the impact of multi-objective compared to single-objective genetic algorithms. In order to be able to re-use problem formulations, we use a Python-implemented interface called *OptProb*, originally devised by Seifert (2007) and only briefly summarized here. *OptProb* does not itself provide any functionality other than data structures and activities harmonizing the interfaces of different optimizers. It can hence be considered an adapter framework.

The main components of *OptProb* are as follows:

- Parameters are defined using the *Parameter* adapter, which provides attributes such as granularity, bound constraints and an initial value. Most of these attributes are optimizer-specific. Only those attributes required by the targeted approach have to be supplied.
- Similarly, objectives can be specified through the *Objective* adapter. Objectives may contain merit functions or solely serve to collect pre-computed results. In the multi-objective case, each criterion is specified by one *Objective* instance. In order for a dual problem formulation, both as a single-objective and as a multi-objective one, an additional aggregate objective can be defined, called *CombinedObjective*. There, additional weight parameters to control the aggregation, which may for example have the form of a weighted sum, can be provided or dynamically assigned by the optimization routine.
- The *Constraint* adapter allows the definition of constraints, generally in combination with a function. Constraints can be classified as linear, non-linear, equality and inequality constraints through attributes.
- The main interface class is called *Problem*. It serves as a container for the former entities but also provides a number of function interfaces, for example, the *evaluate()* routine, called by the optimization algorithm and intended to perform time-consuming computations, which may be deferred to external programs or libraries.

All optimization routines used in this work are either designed or modified to conform to the *OptProb* interface, thus facilitating a straightforward transition from one optimizer to another, provided the problem has also been formulated using *OptProb*.

In addition to the above core functionality, *OptProb* provides a number of convenience routines, for example, to store results for statistic purposes and to visualize optimality sets in order to facilitate the post-optimization decision making process. Figure 6.2 shows an example obtained from the simultaneous source/mask/projector optimization discussed in Section 7.4. On the left (a), the Pareto front (green circles) of two criteria of the multi-objective problem—sidelobe printability minimization and NILS maximization—is shown. Three candidates are manually selected (red circles), one of which is solution 1269, also shown in the spread sheet window (b). For further investigations, the solution can be re-evaluated.



(a)

idx	O: nils	O: cddiff	c: sidelobe	O: pw	P: source1	P: source2	P: source3	P: source4
1260	42.9323989	-1.0602209	0.56218522	0.01655235	942152931	165091670	580823563	907274334
1261	43.1926744	-1.0104240	0.53789405	0.02159933	119435297	985721877	580963132	149823104
1262	42.4702803	-0.8350390	0.57211621	0.01362792	942152931	165091670	580823562	907274334
1263	43.6076828	-1.4790890	0.55146161	0.01848759	942152928	120224999	580963132	149679905
1264	43.8911164	-4.2585971	0.61026721	0.01189748	104300541	165089135	259857686	149682372
1265	45.7777669	-2.8677604	0.53386605	0.00306839	848478045	971873281	616851191	149679904
1266	43.7225846	-1.7107308	0.54474658	0.01982619	104300510	971873281	616993097	907728219
1267	42.6351731	-5.5901035	0.62786590	0.00277476	942152932	971873284	616991929	149679905
1268	46.3484820	-4.5315154	0.56324426	0.01077138	136729272	625966829	616991928	102121893
1269	46.3903310	-2.6201646	0.45071254	0.02945601	104396265	165088921	579274282	149679905
1270	38.4654821	-6.2631484	0.64117277	-0.1173897	906124244	566186915	579558196	149679905
1271	44.5044398	-2.6519857	0.54730003	0.02007364	108895642	971873281	616852291	149679904
1272	47.8787910	-2.3005679	0.43612340	0.03169367	108895642	971873281	616993028	149679909
1273	47.4052871	-3.2563395	0.49142617	0.02107475	104396254	971873281	580963132	149682372

(b)

Figure 6.2: OptProb result browser: (a) Plot of dimension-reduced Pareto front given by two criteria of the multi-objective problem. User-selected individuals are marked red. (b) Spread sheet view of the same population, scrolled to show one of the selected solutions, which could be re-evaluated for further investigations.

OptProb can also be used to encapsulate simulation and evaluation routines. For example, the lithographic process and the respective assessment routines of the mask absorber optimization examples in

Chapter 8 have been combined into a class infrastructure that heavily utilizes the *OptProb* interface. This is schematically shown in Figure 6.3.

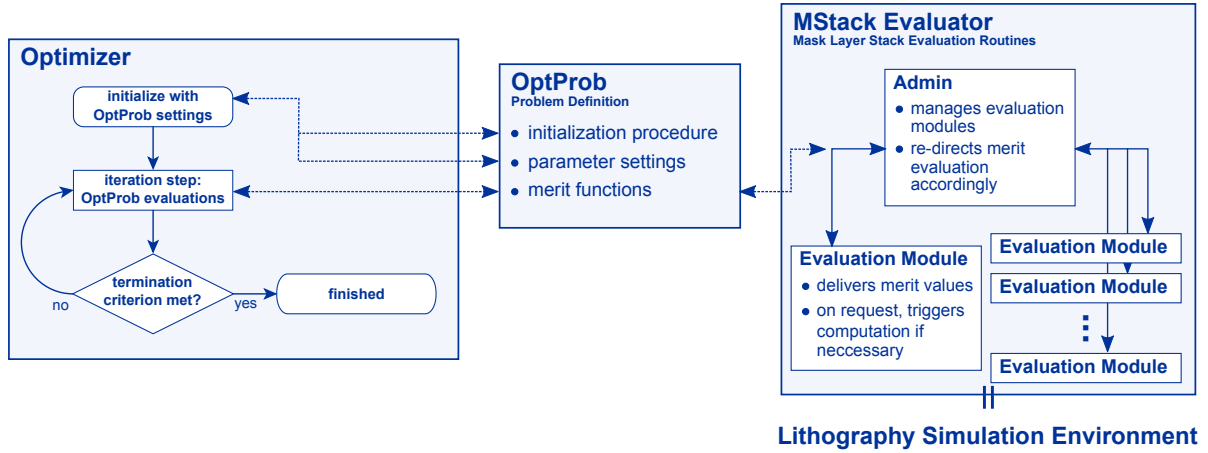


Figure 6.3: Optimization framework, consisting of an optimizer, for example a genetic algorithm, the problem formulation interface (*OptProb*), the evaluation abstraction layer (*MStackEvaluator*) and the simulation environment (e.g., (*Dr.LiTHO*)).

6.5 Task Distribution Environment

Population-based algorithms such as evolutionary or memetic algorithms are well suited to be parallelized. This is because many individual function evaluations are mutually independent and can hence be conducted concurrently. There are different scopes or levels at which an explicit parallelism can be introduced (Cantú-Paz, 1997; Talbi et al., 2008). For example, the subpopulations of island models (Grosso, 1985; Cohoon et al., 1987), which, broadly put, are isolated, individually evolved and then intermigrated,² often serve as an adequate unit of concurrency (e.g., Gustafson and Burke, 2006). This technique is especially well suited for evaluation functions with low runtimes, where a communication overhead presents a significant efficiency penalty. In our case, the opposite is true: Function evaluations can be expected to be long-running, and a moderate communication overhead can be considered negligible. A straightforward master-slave parallelization—as it is called by Cantú-Paz (1997)—is thus suited. For that purpose, we have adopted two alternative mechanisms:

- (1) *Message passing interface (MPI)*: a communication framework which provides a high-level interface for parallelization tasks on distributed memory systems (Gropp et al., 1999). It is available for many computer architectures. Programming bindings include C/C++ and Fortran. Although some implementations (especially of the new version MPI 2) provide mechanisms for fault-tolerance and load-balancing, including these features is intricate and has to be specially tailored per application. In addition, there is a number of different MPI implementations, many of which are not compatible to each other, making it difficult to provide portable programs.
- (2) *A distributing Python tasks environment (DisPyTE)*: an event-driven distribution environment with a certain fault-tolerance and load balancing mechanisms. It is briefly explained in the following.

DisPyTE was developed by Popp (2005) for the distribution of concurrent computational tasks, for example, in the scope of parameter studies or optimization.³

It can be described as a requester/worker framework, where remote worker processes provide arbitrary computation services, and a requester dispatches distributable tasks by conferring them on a broker.

²For example in order to preserve diversity, similarly to niching techniques (Section 4.9).

³Also see Fühner et al. (2006).

As soon as the broker finds a suited, idle node, it assigns it a new dispatched task. Once a worker has finished a task, it sends its result back to the broker. This result is then made available to the requester by triggering a callback function. Figure 6.4 illustrates the main components and principle call scheme of DisPyTE.

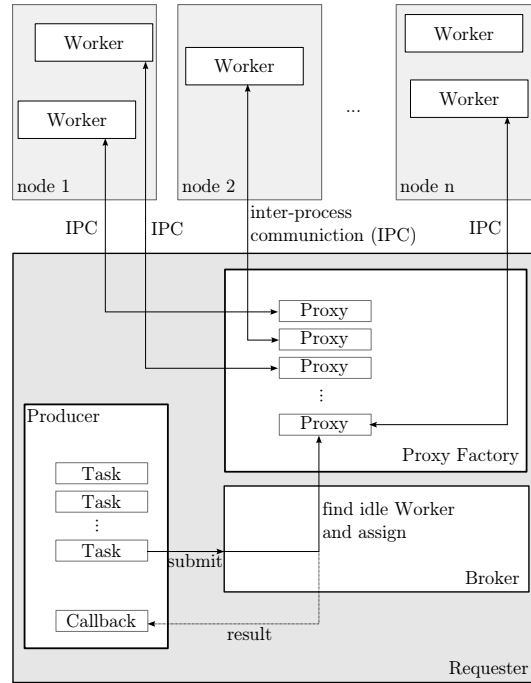


Figure 6.4: Main components and call scheme of DisPyTE.

An *implicit* load-balancing mechanism is inherent to DisPyTE, since the broker distributes tasks only to idle workers. Thus for example, a worker that is twice as fast as another one will in general be assigned twice as many tasks. Once all tasks are assigned, the broker starts to redundantly distribute unfinished jobs. By so doing, very slow nodes are prevented from thwarting the overall job performance. As soon as a job has been completed, the corresponding computation on the redundant worker nodes is canceled.

The proposed approach is also fault-tolerant in that it is not dependent on the success of all nodes. If a worker fails, the pending task is re-assigned to a different worker. Even in case no workers are currently available, the broker will not fail but wait for new computing entities to attach.

DisPyTE does not depend on a specific interprocess communication mechanism. But as a first implementation, the Python package Twisted Matrix proved an ideal framework (Fettig, 2005). It provides mechanisms for an asynchronous, event-driven network communication, directly reflecting the callback mechanisms of DisPyTE. As demonstrated in Figure 6.5, in order to remotely execute a routine, a so-called *deferred* is called. Both a callback and an error callback (*errback*) are attached to this *deferred*. Once the remote process has set a value on the *deferred*, the callback is triggered. In case of an error, the *errback* is executed instead.

One major advantage of the proposed concept is its modularity: Any component can be easily replaced by a custom implementation. Thus an integration of problems defined through the *OptProb* interface is straightforward (see Section 6.4).




Figure 6.5: Twisted *deferreds*: once the remote process sets a value on the *deferred* a previously attached callback is triggered; in case of an error an optional *errback* is called.

6.6 Brokered Network for Hybrid Optimization Approaches

A similar framework as DisPyTE, but specifically tailored to facilitate the hybridization of two or more optimization routines was developed by Dürr (2006a); Dürr et al. (2007a). A related approach was developed by Cahon et al. (2004), albeit their comprehensive ParidesEO framework is better described as a C++ optimization building block library, providing a large number of components for heuristic and deterministic optimizers, evaluation and parallelization (including MPI and grid-computing facilities). A similar Java-based building block framework, not inherently providing parallelization functionality, called MAFRA (Memetic Algorithm FRAmework) was devised by Krasnogor and Smith (2000).

Heavily making use of communication agents, termed brokers, the environment used in this work has the recursive acronym “LisBON is brokered optimization network (LisBON).” Like DisPyTE, it makes use of an event-driven interprocess communication, allowing for function evaluation workers but also optimization processes to be distributed onto multiple machines. In the current implementation, which is exclusively realized in the programming language Python, an XML-RPC⁴ protocol is used for the underlying interprocess communication. In the current version, LisBON does not employ any fault tolerance. It is hence best suited to be executed in dedicated cluster environments or multi-processor machines.

In order for LisBON to be applicable, two prerequisites have to be met:

- (1) The employed optimization algorithms have to be separable into individual steps. A step may for example consist of one or more GA generations or iterations of a local search routine. Most algorithms, such as the GA, the SQP or the NM methods either inherently fulfill this requirement or can be adapted such.
- (2) The algorithms have to be designed such that they operate on one or a set (population) of solution candidates, which is the case for all methods considered in this work.

broker The broker constitutes the central entity as it controls the overall flow of the hybridized, distributed system. Any optimization step or function evaluation is instantiated by the broker. The broker thus also manages the communication and data exchange between the workers. Additionally, when supplemented with respective rules, the broker can spawn new processes, including function evaluators or search algorithms.

rule Rules determine the general flow of the hybrid optimization system and are hence processed by the broker. They can be considered small control code pieces triggering an action. As an example, also the trigger and selection mechanisms explained in Chapter 5 have been realized through the

⁴A remote procedure call (RPC) mechanism using extensible mark-up language (XML) as the protocol language container and standard transport layers such as hypertext transfer protocol (HTTP).

definition of such rules. Other rules, for instance, governing the transfer of population containers or implementing a (self-)adaptive selection of memes could also be implemented.

worker Similarly to DisPyTE, LisBON introduces the notion of workers, which in contrast, however, are not limited to function evaluation entities but also encapsulate the optimization algorithms to be hybridized. This concept allows, for example, the concurrent execution of several local search instances.

population container Candidate solutions are transferred from and to the broker by means of containers that encompass a single solution or a population of solutions but that can also be annotated, for example, with additional state information optimization algorithms may depend upon.

Both the memetic algorithm (Section 5.3) and the memetic algorithm with function approximations (Section 5.6) used in this work are realized using LisBON. This is greatly facilitated by the fact that the LisBON framework complies with the *OptProb* interface (see Section 6.4).

7 Source/Mask Optimization

As shown before, lithographic process conditions are determined by a number of interacting properties, related to the illuminator, the projection system and the wafer stack, all of which can be individually optimized. Many of these actuating factors, however, mutually influence each other such that an isolated tuning of just a small set of parameters may easily lead to a suboptimal total lithographic performance. On the other hand, to render an optimization effort feasible and tractable it is inevitable to identify and separate units of the complete system, which then can be dealt with in a rigorous manner. One of such units is the lithographic imaging performance to whose improvement this chapter is devoted. As discussed in Chapter 2, without the utilization of advanced resolution enhancement techniques, the current optical lithography generation—operating at a wavelength of 193 nm—would fail to meet the requirements for state-of-the-art semiconductor technology nodes. Indeed, an enormous amount of efforts have been put forth to exceed the given physical limits. The obtained techniques attack all aspects of the lithographic system. And while many of these measures are now well-established either in lithography or in the scope of other optical instruments, their combination still poses a significant difficulty. It is for that reason that in a field that has long been utilizing all types of computer programs, the term *computational lithography* becomes increasingly popular. It also draws a future perspective, in which the set of available techniques such as RET is considered the toolbox from which an automated procedure assembles an appropriate combinations. Our effort, a rigorous source/mask optimization approach, is directed toward that goal.

As will be shown, in many of the related studies the approach is termed inverse lithography technique (ILT). Inverse problems in the scope of imaging problems can be broadly described as follows (Saleh and Sayegh, 1981): Given a prescribed output pattern, (an image or a resist pattern), a known forward transformation (e.g., imaging and resist response), the goal is to find an input pattern (mask) under which the prescribed output is obtained. For this problem to be well-posed, the forward transformation would have to be invertible. It is not, however, even without any regard to the source. To illustrate that the problem is highly ill-posed, just consider the band-limiting property of the projector, due to which numerous differently shaped masks lead to the same diffraction pattern. Thus a more relaxed formulation for the problem has to be established, such as a minimization of the error between prescribed and obtained output pattern. Yet, even such a formulation is highly ambiguous since any optimum will be strictly local. In other words, the search space is non-convex but instead highly multimodal. Consequently, it is often proposed to extend the problem by additional constraints in order to regularize it, that is, to reduce the degree of multimodality.

We believe that the problem is not an ill-posed inverse problem but ill-stated as an inverse problem, mainly for the following reasons.

- The aim is not to obtain a prescribed pattern. Rather, the output pattern should be an optimum with respect to or in some cases be constrained by a number of criteria, such as a minimum edge placement error, a high contrast and a stable through-focus behavior etc. Minimization of the difference between a candidate and a target image strongly biases the approach toward that target solution. The problem is then to generate a target such that this bias actually conforms to the intention.
- The aim is not to obtain just any input pattern. The regularization efforts clearly demonstrate that. Rather, the mask should meet a variety of manufacturability and precision criteria (such as a minimum MEEF).
- Process variability is prohibitive for a concise formulation of the problem as an inverse problem, hinted upon by the many robustness and sensitivity criteria incorporated into the approaches.

To avoid any misunderstanding, a large number of the approaches, also the ones termed ILT, do account for the problems highlighted above. But by doing so, they in fact do not treat the problem as an inverse formulation but as it is, a plain optimization problem. We will hence stick to the term source/mask optimization (SMO).¹ There is one exception, when, as proposed by some, the problem is separated into two parts: (1) the search for an optimum wavefront and (2) the search for an ideal mask corresponding to that front. The second subproblem could indeed be termed “inverse.” This property, however, can also be considered one of the greatest issues of this approach: At worst, any manufacturability constraints may foil the previous optimization step since the regularized mask may yield a diffraction spectrum rendering the total solution suboptimal.

Before we present the details and some results of our approach, we shall review the field of related work.

7.1 Related Work

Few of the ideas incorporated in SMO are entirely new. In fact, many of the techniques, especially those associated with mask optimization, have been either developed for or are already used in conventional OPC software (see Section 2.6.1). Two aspects distinguish SMO from traditional approaches: (1) The rigorous optimization of the source, which before was adjusted empirically or heuristically, and (2) the awareness that sources and masks have to be co-optimized in order for both to constructively combine to a stringent wavefront control.

Early rigorous optimization approaches include a series of papers in which the image synthesis problems for different illumination conditions were considered (Saleh and Sayegh, 1981; Saleh and Nashold, 1985; Nashold and Saleh, 1985; Sherif et al., 1995). There, the problem was stated as a search for an input pattern (mask) that minimizes the error between a prescribed image and the image under the input pattern. These approaches already accounted for a hard-limiting detector such as a photoresist. Different optimization routines including linear programming or mixed integer programming approaches were applied. A related approach was proposed by Liu and Zakhor (1992). By subdividing the mask into regions for the transition from dark to bright areas (and vice versa) and by maximizing the image slope at the corresponding edges, their approach was able to yield solutions with a significantly improved contrast. For an approach suited in the regime of partial coherent imaging, the authors proposed the use of simulated annealing (see Section 4.1). Using their optimal coherent approximation, commonly termed SOCS (see Section 3.2.6) Pati and Kailath (1994) devised an optimization strategy similar to the one proposed by Saleh and Nashold (1985), albeit accounting for partial coherence.

Oh et al. (1999) proposed a stochastic source/mask optimization routine based on a pixelated representation of the mask and a parametric representation of stepper parameters. The source optimization is performed subsequently. The goal is to find an aerial image that closely resembles the target pattern. The mask optimization approach is similar to the procedure proposed by Saleh and Sayegh (1981): With each iterate, a pixel of the pixelated mask is randomly chosen and inverted. The pixel flip is accepted if and only if it reduces the resulted error. A two-fold manufacturability criterion is included: (1) During the optimization process, clusters of pixels are preferred over isolated ones. (2) A post-optimization routine regulates the mask such that it shows long edges and a reduced number of staircases. During the illumination optimization, the exposure latitude for a set of user-provided sample points is maximized. Due to the limited number of parameters, an exhaustive search is performed. The feasibility of the approach was shown for a number of cases, including chrome-on-glass (COG) masks and phase-shifting masks (PSMs). The authors emphasize the capability of their approach to automatically place serifs. A similar pixel-flipping procedure, used by the former author for the mask only, was later also proposed by Kim et al. (2003).

The rest of this section is devoted to the more recent developments, which will be discussed in more detail.

¹This term has been first used by Rosenbluth et al. (2002).

7.1.1 Wavefront-based Source/Mask Optimization

Several authors have proposed to employ RET optimization approaches by rigorously accounting for their effect in the wavefront domain. Accordingly, customized sources were, for example, proposed by Burkhardt et al. (1998), Smith and Zavyalova (1998) and Gau (2000), and Smith (2002) presented a seminal study on the joint exploitation of custom sources, PSMs and SRAFs. One of the first mutual source/mask optimization frameworks was devised by Rosenbluth et al. (2002). Their technique can be considered as a hybrid local/global search approach: A global optimization approach is employed on a simplified representation of the problem to generate initial solutions for a local optimization step in which a refined version of the problem, including additional figures of merit and constraints, is used. Rosenbluth et al. point out that previous approaches were concerned with the optimization of either the source or the mask, but that an effective wavefront engineering depends on the interaction of both. In other words, the quality of a mask, in terms of image performance, strongly depends on the source for which it has been designed, and vice versa. One approach to this notion is to perform mask and source optimizations in separate disjoint steps, for example, by first optimizing the source and conducting a mask optimization ex post. This approach, however, necessarily leads to sub-optimal solutions since the second optimization step faces prematurely fixed conditions. Rosenbluth et al. hence devised a more involved procedure incorporating multiple optimization stages, as schematically illustrated in Algorithm 7.1.²

Algorithm 7.1: Multi-stage Optimization Strategy

```

1  $S' \leftarrow 0$  /* initialize minimum pupil fill */
2  $\mathbf{s} \leftarrow$  initialize source variables
  /* global optimization */
3  $\mathbf{a}^* \leftarrow$  globally optimize mask diffraction pattern
4 while termination criterion not met do
5    $\mathbf{s}^* \leftarrow$  globally optimize source variables for given  $\mathbf{a}^*$  and a given minimum pupil fill  $S'$ 
6    $\mathbf{s}^*, \mathbf{a}^* \leftarrow$  locally refine  $\mathbf{s}^*, \mathbf{a}^*$ 
7    $S' \leftarrow S' + \Delta S'$  /* increase pupil fill level */
  /* local refinement */
8  $\mathbf{s}^*, \mathbf{a}^* \leftarrow$  optional local optimization of  $\mathbf{s}^*, \mathbf{a}^*$  with additional criteria (e.g., DOF) for obtained  $S'$ 
  /* mask computation */
9  $\mathbf{T} \leftarrow$  generate mask pattern for  $\mathbf{a}^*$ 

```

Similarly as in previous studies, for example, by Burkhardt et al. (1998), Rosenbluth et al. use a source representation that is closely related to the target layout. Since periodicity is common property in lithographic applications, at least in the frequency-centric treatment normally used in computational lithography, a grid representation can be employed for the diffraction order analysis in the pupil plane. Given a layout as shown in Figure 7.1a, with a pitch $\mathbf{a}_1 := [p_x, p_y]^T$. In the spatial domain of the layout, the pitches between arbitrary features can be written as follows:

$$\mathbf{p} = m\mathbf{a}_1 + n\mathbf{a}_2 \quad m, n \in \mathbb{Z}, \quad (7.1)$$

where \mathbf{a}_1 is defined as before and

$$\mathbf{a}_2 := [-p_x, p_y]^T.$$

Comparably to the reciprocal lattice in solid state physics and following Burkhardt et al. (1998), one can find pupil plane equivalents (\mathbf{g}_1 and \mathbf{g}_2) for the two vectors, satisfying the following condition

$$\mathbf{a}_i \mathbf{g}_j = \begin{cases} 2\pi & i = j \\ 0 & i \neq j \end{cases} \quad i, j = 1, 2. \quad (7.2)$$

²Loosely adopted from Rosenbluth et al. (2002).

This can be motivated by noting that for the Fourier transform of the mask transmission function, which is equivalent to the Fraunhofer diffraction formula, we have that

$$\int T(\mathbf{r})e^{i\mathbf{k}\mathbf{r}}d\mathbf{r}, \quad (7.3)$$

where \mathbf{r} is the spatial coordinate and \mathbf{k} is the pupil plane coordinate. The mask transmission function can then be defined as the following Fourier series

$$T(\mathbf{r}) = \sum_{m,n} \rho_{m,n} e^{i(m\mathbf{g}_1 + n\mathbf{g}_2)\mathbf{r}} \quad m, n \in \mathbb{Z}. \quad (7.4)$$

This Fourier analysis leads to a straightforward determination of the diffraction orders in the pupil plane: The integers m and n represent the number of the orders, whose position in terms of the normalized spatial frequency in the pupil plane is given by a linear combination of the vectors \mathbf{g}_1 and \mathbf{g}_2 . This is depicted in the diagram shown in Figure 7.1b, which was obtained for a $[187.5, 75.]^T$ nm grid, an NA of 0.8 and a wavelength of 193 nm. About each diffraction order, a circle with a radius corresponding to the NA is centered, leading to a number of overlapping segments, called archels. Each archel then defines a region in the source that directs a specific combination of diffraction orders into the projector lens. These archels, more precisely their intensity, and potentially also their phase and polarization, present the optimization variables for the source.

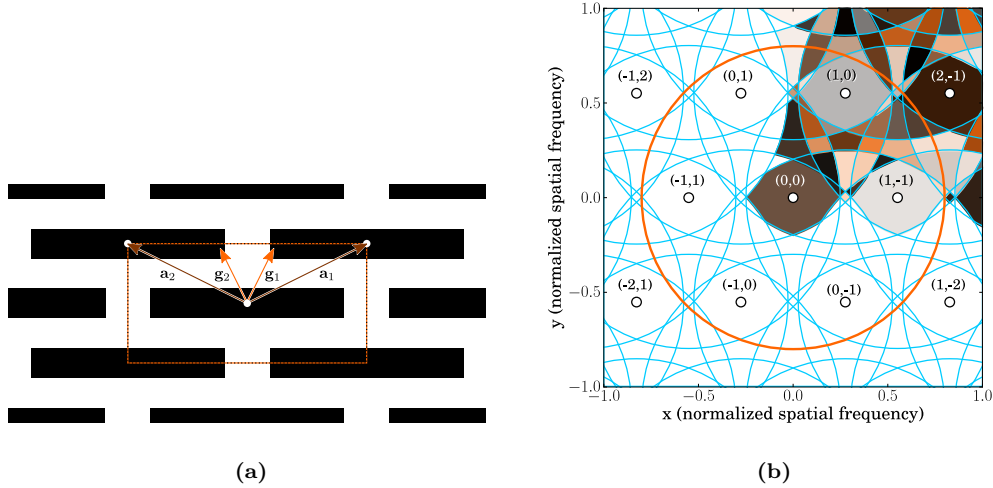


Figure 7.1: Archel-based source representation: Typically computational lithography employs periodic structures or boundary conditions such as the layout in inset (a). The two-dimensional period can then be given as vectors \mathbf{a}_1 and \mathbf{a}_2 . The corresponding basis vectors in the pupil plane are \mathbf{g}_1 and \mathbf{g}_2 . Any linear combination with integer scalars determines a diffraction order in the coherent (axial) illumination case (b). By centering circles with radius equal to the NA about each diffraction order, one can construct unique source segments that direct the same orders into the pupil—coherent on-axis illumination assumed, again. Each of the so-obtained segments or archels can hence be used as one optimization variable. The staggered pitch (\mathbf{a}_1) in this example is $[187.5, 75.]^T$ nm, the NA (black circle) is 0.8 and a wavelength of 193 nm is assumed (adopted from Rosenbluth et al., 2002).

Similarly, the mask response is defined as its amplitude per diffraction order. This technique is sometimes referred to as wavefront engineering: Instead of varying the mask layout and optimizing it for a desired image behavior, first the wavefront response of the mask is optimized for, and the so-obtained continuous map is then reproduced by a suited transmission function in a second step. In

addition to a concise mathematical formulation and a complexity reduction of the first step, the major advantage of this is its independence from a specific mask technology. As a disadvantage, however, mask manufacturing constraints can be employed only in the second step. Thus compromise solutions, highly manufacturable with an acceptable imaging behavior, can be considered only insofar as they are covered by the optimized continuous diffraction amplitude function.

In the first step of the global optimization stage (Algorithm 7.1, l. 3), a simplified replacement problem is first posed. This is done by considering only one source region at a time, for which the mask diffraction amplitudes are optimized. The goal of this optimization step is the minimization of (sampled) dark areas while at the same time ensuring that samples of the bright regions actually exhibit an intensity of one and above.³ Points on feature edges can be emphasized by an introduction of weights. As pointed out by Rosenbluth et al. (2002), the so-stated quadratic problem with quadratic inequality constraints is highly multimodal. In order to obtain a globally (quasi) optimal solution, Rosenbluth et al. propose a decomposition they term joint eigenvector formulation, in which a common eigenspace for the dark and the bright region imaging matrices is set-up (which due to homogeneity of both matrices can be done through a coordinate transformation). Solutions are assumed to consist of a combination of the principal eigenvectors (Rosenbluth et al. (2002) use three). This is so because each bright point must fall into a bright region of the eigenvector, but a number of eigenvectors do not capture the intended intensity distribution and thus degrade the contrast. The minimization problem can then be regarded as a search for the closest point to the origin of the eigenspace, which does not violate the bright constraints, i.e., which does not lie within any of the ellipsoids spanned by the constraints. The search domain is decomposed by introducing a celestial hyper-sphere, representing an abysmally high intensity. A triangular decomposition of the search space is produced by first projecting the principal axis of the constraint-defining ellipsoids onto the sphere. Additional nodes on the sphere are then generated through subdivision. Finally, solutions are obtained by searching for the closest intersection of the edge between the vertex on the sphere and the origin and a subsequent local search in the vicinity of this intersection, yielding an enumeration of different locally optimal solutions. The global solution can, for example, be taken as the best of the so-obtained local candidates.

This solution is used in the following source optimization step (ll. 4–7), for which Rosenbluth et al. (2002) employ a linear program formulation. A mutual local refinement of the source and the mask diffraction variables can be optionally employed (l. 6). Due to the non-linearity of this problem, sophisticated solvers, for example, using a SQP formulation (see Section 5.4.1), suggest themselves. The source optimization steps are repeated for different pupil fill levels (S'). Subsequent to this iteration, a compromise solution balancing an ideal source/mask combination and an acceptable pupil fill can be selected, either manually or automatically. In a final step (l. 8), the so-obtained solution is again locally refined, optionally using a more rigorous problem formulation. For example, while in the former steps only the best-focus performance is considered, the local refinement step may incorporate image conditions for several focus positions (Rosenbluth et al., 2002).

In the last step (l. 9), the mask diffraction variables are translated into a pixelated representation of the mask layout. Like the other steps, this is an ill-posed problem, since a number of different mask functions generally conform to the optimized amplitude distribution. It is thus posed as an optimization problem. Using a pixelated mask representation, Rosenbluth et al. (2002) propose to maximize the total intensity of the gridded mask, subject to the optimized diffraction amplitude distribution. As a simple manufacturing criterion, a pixel size conforming to the mask writer resolution is used. The criteria for the mask computation of course depend on the employed mask technology. Especially in the case of alternating phase shifting masks, design conflicts have to be resolved. The authors hence introduce a final regularization step, in which design conformance and other mask-related criteria can be taken into account.

One shortcoming of the original optimization approach is the restriction of quality assessment to the best focus position. To circumvent this issue, Rosenbluth et al. (2002) conducted an experiment in which the DOF performance is optimized for only in the final local refinement optimization step and were able to compute well performing source/mask combinations with this modification. In a later

³Intensities in this step are relative.

paper, Rosenbluth et al. (2007) indicate that also the global optimization step could be conducted at the boundaries of a pre-defined DOF range.

Other extensions include the optimization for process window performance during the source optimization stage (Rosenbluth, 2006), the consideration of thin-film effects and the incorporation of circuit-design-driven criteria (Rosenbluth et al., 2007).

7.1.2 Interference Map Lithography Approach

Socha et al. (2005) propose an aerial image-based source/mask optimization approach, in which, similarly to Rosenbluth et al. (2002), first an ideal mask spectrum is sought and in a second step a corresponding mask transmission function for this spectrum is generated. Employing an Abbe imaging approach, the impact of each point of a pixelated source on the intensity distribution is individually evaluated, for example, on representative clips in the image domain. To optimize for image fidelity, Socha proposes to measure the ILS at critical edges. The optimization parameters of the mask are given by the discretized mask spectrum, restricted to the diffraction orders collected by the projector. Since the overall goal of the optimization procedure is the maximization of both the exposure latitude (EL) and the DOF, the minimum ILS for the representative edges is maximized over a pre-defined focus range. The minimax problem posed so is additionally subject to a number of constraints, including safeguards to ensure that the intensity in dark and bright fields is within an intended range and illumination and mask spectrum regularization constraints. This renders the optimization problem generally non-linear (cf. Section 5.4.1). In addition to high dimensionality, the problem exhibits a high degree of multimodality (cf. Rosenbluth et al., 2002). Subsequent to the co-optimization, the spatial correspondent of the mask spectrum (the pseudo-near field) determines the mask layout. There, the continuous field has to be translated into a manufacturable geometry consisting of regions with discrete transmission values.

In the same paper, Socha also proposes a source/mask procedure suited for larger areas, combining the former rigorous approach for critical structures and an interference mapping lithography (IML) approach for an assist feature (AF) placement for the remaining features. Since in this second phase the source can be assumed to be constant, the SOCS approach can be applied (see Section 3.2.6). The computation of the interference map is loosely related to a principal component analysis: Socha et al. (2004) argue that in many cases only few SOCS kernels exhibit a significant eigenvalue. Specifically the first kernel carries most of the information of the illumination/projection system. It can hence be regarded as an approximate analogon of the point spread function (PSF) for partially coherent illumination. The interference map is computed by convolving the first SOCS kernel with the intended layout,⁴ approximated through a linear combination of Dirac delta and line functions. Its discretized version is then used to optimize the exposure latitude by maximizing the intensity in bright regions and minimizing it in dark regions. Similarly, the depth-of-focus can be optimized for by minimizing the intensity change with respect to a predefined focus range. Socha points out that the two criteria, EL and DOF, are incommensurable. As a solution, he proposes to reformulate the problem such that the EL component becomes a constraint. As an alternative, it could be stated as a bi-objective problem and accordingly solved with a multi-objective optimization approach. Socha et al. (2005) devise another technique to optimize the source, broadly put, by maximizing the cross-correlation of the first SOCS kernel and the pitch frequency, which is a measure of the occurrence of pitches in the target pattern.

7.1.3 Pixelated Source Optimization

Granik (2004) developed a source optimization aimed at maximizing the throughput of illuminators using DOEs. Mainly to account for their discrete nature, he also employs a pixel-based source representation for the optimization. The pixel grid is defined on a Cartesian coordinate system, bounded by a circle of radius σ . The equidistant $N \times N$ ($N \in \mathbb{N}, N$ odd) grid is hence given by:

$$X := \{(x_i, y_j) | (i, j \in \mathbb{N}, i, j < N) \wedge ((i - N/2)\Delta x)^2 + ((j - N/2)\Delta y)^2 \leq \sigma^2\}$$

⁴Socha also indicates how to extend this approach to account for additional kernels.

with $\Delta x = \Delta y := 2\sigma/N$. The source function maps the grid to its intensity values:

$$S : X \rightarrow \mathbb{R}_0^+.$$

In the regime of DOEs, light is not blocked, but the total energy E_0 emitted by the illuminator is transmitted. Formally

$$E_0 = \sum_{i=0}^N \sum_{j=0}^N S(x_i, y_j).$$

One of the optimization criterion is to minimize intensity peaks in individual pixels (or DOEs). This prevents damages in lenses and especially their coatings:

$$\min \max S(x_i, y_j). \quad (7.5)$$

One solution and the lower bound to this criterion, is a uniform intensity distribution:

$$S(x_i, y_j) = E_0/N^2 \quad (\forall i, j \in N).$$

Granik proposes to replace the harsh and intricate minimization criterion of (7.5) with a Sobolev norm, which imposes a smooth distribution by including derivatives:

$$|S|_{2,2} = \sqrt{\sum_{l=0}^2 \sum_{i=0}^N \sum_{j=0}^N (\alpha_l S^{(l)}(x_i, y_j))^2}, \quad (7.6)$$

where α_l are metric constants to control the degree of smoothness.

As the central image optimization goal, Granik introduces what he calls the image fidelity error:

$$F := |I - I^*|_2, \quad (7.7)$$

where I' denotes the target image. Granik argues that this metric is capable of capturing local image properties, in contrast, for example, to global metrics such as the global contrast (which is adequate only for grating-like structures). But the determination of an appropriate target image is far from being straightforward and bears the risk of a strong bias toward a suboptimal solution, predetermined by the target.

Granik (2004) exemplifies the use of a target image by assuming that the target is equivalent to the non-projected mask function. By substituting Hopkins' formula for the intensity calculation in (7.7), we obtain:

$$\begin{aligned} (\mathcal{F}(I) - \mathcal{F}(I^*))(f, g) &= \iint_{-\infty}^{\infty} (\mathcal{T}((f' + f, g' + g), (f', g')) - 1) \\ &\quad \mathcal{F}(O)(f' + f, g' + g) \mathcal{F}(O)^*(f', g') df' dg'. \end{aligned} \quad (7.8)$$

Thus in order to minimize the error, the TCC components (\mathcal{T}) should be as close to one as possible.⁵

To intensify the optimization for critical layers or aspects of the image, Granik proposes the use of a weighting function. Or—also to speed up the optimization—a reduction to one-dimensional cut-lines; thus reducing the error evaluation to multiple one-dimensional functions:

$$F_j = |I_j(t) - I'_j(t)|_2, \quad (7.9)$$

where j denotes the (index of the) cut-line and t the position along this cut-line. Since the Euclidean norm can be considered to yield the average distance between the real and the target image, extreme

⁵ As Granik points out, this consideration may motivate an archel-based source representation, using archels corresponding to non-zero kernels (see Section 3.2.6) or even a truncated number of significant kernels (cf. Section 7.1.1, Section 7.1.2 and Burkhardt et al. (1998)).

variances may be underestimated. In contrast, the uniform norm (or supremum norm) measures the maximum distance between the images and is hence well-suited to capture large deviations, which might jeopardize the printability of the entire layout. However, this norm poses a number of issues for the optimization problem, most importantly, its non-transferability into the frequency domain.

Similarly to the source smoothness problem, Granik makes use of the $W^{1,2}$ Sobolev norm:

$$F_{W^{1,2}} = \alpha_0^2 \|I - I'\|_2^2 + \alpha_1^2 \|(I - I')^{(1)}\|_2^2, \quad (7.10)$$

where α_0 and α_1 define the ratio of the distance measure between the intensity values and their derivatives. In the vicinity of boundaries of features, the first derivative corresponds to the image slope. Since this slope is most likely infinite in the target image, this error definition can be used to implicitly maximize the image slope, as illustrated in Figure 7.2. Similarly to a more explicit evaluation of the image slope such as the normalized image slope (NILS), introducing such a criterion yields solutions with superior dose latiti




Figure 7.2: Different slopes in target and actual pattern at feature edges. While the target pattern is typically defined to have an infinite slope, the actual pattern may exhibit a significantly different gradient. By introducing a derivative-aware norm, such as the Sobolev norm, the slope difference can be implicitly introduced as an optimization criterion.

In his optimization routine, Granik also accounts for a stable through-focus behavior. For that, fidelity error evaluations for different focus positions are introduced. Two alternatives are proposed: (1) an optionally weighted average over all focus positions, or (2) an image evaluation only at one defocus position. Granik argues that the latter significantly reduces the run-time, but results are biased towards the pre-selected focus position. Thus a careful, problem-dependent choice of one of the two alternatives has to be carried out. The optimization problem can then, for example, be stated as the following bi-objective form, aiming at finding a source function that both minimizes the image fidelity error and that yields an ideal intensity distribution:

$$\begin{aligned} & \arg \min_S \quad [F^+, \|S\|_{2,2}]^T \\ & \text{subject to} \quad S \geq 0 \\ & \quad \quad \quad \|S\|_1 = E_0 \\ & \text{where} \quad S \text{ is the source function to optimize} \\ & \quad \quad \quad F^+ \text{ is the image fidelity error at a predefined defocus position.} \end{aligned} \quad (7.11)$$

Equation (7.11) could be straightforwardly solved with a multi-objective approach such as an epsilon-constraint technique (see Section 4.11.1), which would hence allow for an ex posterior selection of suited compromise solutions. Granik (2004), however, poses the problem as a weighted least squares replacement problem. There, the source intensity is replaced by a relative source intensity. In other words, the intensity is normalized by the energy:

$$j := \frac{S}{E_0}.$$

Moreover, Granik uses a penalty method (see Section 4.4) to account for the energy constraint in order to obtain the following concise least-squares formulation:

$$\begin{aligned} \arg \min_j \quad & w^2(F^+)^2 + (1-w)^2|j|_{2,2}^2 + c^2(|j|_1 - 1)^2 \\ \text{subject to} \quad & j \geq 1 \\ \text{where} \quad & w \in [0, 1] \text{ is a weight parameter and} \\ & c \text{ is the penalty coefficient } \geq 1 \end{aligned} \tag{7.12}$$

The minimization term in (7.12) can be considered a function of the source point. It is thus convenient to employ Abbe's approach for the image calculation (see Section 3.2.5), which directly resembles this situation. For a given mask and precomputed coherence terms (a_σ in (3.8)) and assuming a discrete image, the intensity calculation can be stated in the following matrix form:

$$\mathbf{I} = \mathbf{T}\mathbf{j},$$

where the transformation \mathbf{T} is precomputed from (3.8) and \mathbf{j} denotes an appropriate flattened vector of the intensities of the pixelated source. The problem is then solved with the least-squares method.

The results reported by Granik (2004) clearly demonstrate a superior imaging performance of his approach, especially in comparison with standard parametric and contour-based source definitions. Granik showed a strong dependence of the result on the weights assigned to different image regions.

7.1.4 Image Fidelity Mask Optimization

Also for the mask optimization, Granik (2006) proposes a pixel-based approach. The size of each pixel is determined, for example, by the resolution of the mask writer, thus already incorporating one mask manufacturability criterion into the problem definition. Moreover, each pixel can take arbitrary phase and transmission values, allowing for different mask types—including binary and phase-shifting masks—to be taken into account. Like for the source optimization, the figure of merit is given by the difference of the aerial image resulted by the mask of a solution candidate and a target image, which may incorporate the intended binary pattern and the desired image and contour fidelity. In addition, resist characteristics can be incorporated, for example, by using a simple threshold approach. Granik studied three different optimization techniques:

- (1) A line search approach,
- (2) a quadratic programming technique and
- (3) a general non-linear method.

For each, the problem has to be formulated in a slightly different way.

For the linear approach, the Hopkins imaging equation (see Section 3.2.6) is approximated by a linear coherent equation, which is equal to the SOCS formula (3.14) truncated to only the most effective kernel (compare Section 7.1.2), so that we have:

$$\mathbf{I}[l_x, l_y] \approx \lambda |(\mathbf{x} * \mathbf{O})[l_x, l_y]|^2,$$

where λ and \mathbf{x} are the eigenvalue and the eigenvector, respectively, and \mathbf{O} is the mask transmission matrix. For the amplitude, we obtain:

$$\mathbf{A}[l_x, l_y] = \sqrt{\lambda}(\mathbf{x} * \mathbf{O})[l_x, l_y] \tag{7.13}$$

This simplification allows for the application of an iterative approach first proposed by (Saleh and Nashold, 1985): The image is assumed to be an output signal g due to an input signal f , the mask, with

$$g = LBf, \quad (7.14)$$

where B is a bandlimiting operator, forcing the image into the frequency support of the projector, and L is a hard-limiting operator, imposing a threshold cut-off on the result. The output signal g then serves as an input for the next iteration step. Formally,

$$f^{(k+1)} := LBf^{(k)}. \quad (7.15)$$

This procedure is inspired by the Gerchberg-Saxton algorithm (Gerchberg and Saxton, 1972), used to retrieve phase values under known intensity distributions and imaging conditions. This algorithm can be said to be an image reconstruction algorithm in that the object is iteratively modified until close enough an approximation for the target image is obtained. Fienup (1978) proposes an extension to this scheme that relaxes the threshold enforcement operator, which has been shown to lead to a convergence improvement. Granik summarized and parameterized Fienup's modifications as follows:

$$f^{(k+1)} := (LB + \alpha(\gamma(LB - B) - L + I)) f^{(k)}, \quad (7.16)$$

where α and γ are parameters whose choice is critical, which Granik demonstrated by an example in which α was set to zero, which in turn corresponds to a conventional Gerchberg-Saxton approach: Instead of a phase alternation between lines and spaces, a non-useful mask with phases alternating along the line was found. However, with a different choice of operator parameters and a smoothing of the target image, the algorithm converged to a printable solution.

As an alternative, Granik proposed to formulate the difference minimization between a target and the obtained image as a quadratic problem, a formulation closely related to the global mask optimization step proposed by Rosenbluth et al. (2002).⁶ Like in the source optimization case, the optimization goal is hence to find a mask such that the resulted image shows the highest possible agreement with a predefined target image. Mathematically, this can be achieved by finding a mask whose image shows a minimum deviation from the target pattern:

$$\arg \min_{\mathbf{O}} |\mathbf{I}(\mathbf{O}) - \mathbf{T}|_P, \quad (7.17)$$

where \mathbf{T} is the target image and P is the norm used as a measure for the deviation. For a convenient representation, Granik proposes the Manhattan norm, i.e., $P := 1$, which can be interpreted as a linear average.⁷

By additionally assuming a binary target pattern, that is, bright regions have a value of 1 and dark areas are 0, the difference in the discrete case can be stated as follows:

$$\begin{aligned} \varepsilon(\mathbf{O}) &= \sum_{l_x, l_y} |\mathbf{I}[l_x, l_y] - \mathbf{T}[l_x, l_y]| \\ &= \sum_{l_x, l_y} w_{l_x, l_y} (\mathbf{I}[l_x, l_y] - \mathbf{T}[l_x, l_y]) \quad \text{with } w_{l_x, l_y} = \begin{cases} 1 & \mathbf{T}[l_x, l_y] = 0 \\ -1 & \mathbf{T}[l_x, l_y] = 1 \end{cases} \end{aligned}$$

The target pattern is constant and can hence be omitted. Moreover, a weighting function or matrix \mathbf{W} can be introduced to attenuate or ignore the merit of specific pixels or regions:

$$\varepsilon_{\mathbf{W}} : \mathbf{O} \rightarrow \sum_{l_x, l_y} \mathbf{W}[l_x, l_y] \mathbf{I}[l_x, l_y]. \quad (7.18)$$

⁶See Section 7.1.1.

⁷For which Granik points out that it may not be adequate as the performance of the image is only as good as the most poorly performing feature of it.

Substituting (7.18) into the Hopkins imaging equation, integrating the weighting function into the TCC kernel, and applying an appropriate stacking operator allows for a compact, QP compatible representation of the error:

$$\varepsilon_{\mathbf{w}} : \mathbf{O} \rightarrow \mathbf{O}^* \mathbf{Q}_{\mathbf{w}} \mathbf{O}, \quad (7.19)$$

where \mathbf{Q} is the stacked weighted kernel matrix. In order to bound the problem—it can be expected to be non-convex—the box-constraint limiting the mask amplitudes to 1 suggest itself, yielding the following QP:

$$\begin{aligned} \arg \min_{\mathbf{O}} \quad & \mathbf{O}^* \mathbf{Q}_{\mathbf{w}} \mathbf{O} \\ \text{subject to} \quad & g(\mathbf{O}) = |\mathbf{O}|_{\infty} - 1 \leq 0, \end{aligned} \quad (7.20)$$

The problem of (7.20) exhibits a strong resemblance with trust-region models, which are often applied to render non-linear problems amenable for quasi-Newton methods (see, for example, Section 5.4.1). Often the Euclidean norm is employed to define a (iteratively modified) hyper-spherical trust region. By also replacing the maximum by the Euclidean norm in his study, Granik applied a theorem by Moré and Sorensen (1983),⁸ stating that \mathbf{O}' is a global solution to problem (7.20)—albeit with a Euclidean norm—if and only if there exists a non-negative scalar λ , such that

$$\begin{aligned} (\mathbf{Q}_w + \lambda I) \mathbf{O}' &= 0 \\ \lambda(1 - |\mathbf{O}'|) &= 0 \\ (\mathbf{Q}_w + \lambda I) &\text{ is positive semidefinite} \end{aligned} \quad (7.21)$$

Assuming that the matrix \mathbf{Q}_w is indefinite, Granik arrives at the conclusion that the global solution (1) has an active trust-region constraint and (2) is given by the eigenvector corresponding to the smallest negative eigenvalue of \mathbf{Q}_w . The Langrange multiplier in this case is then given by the absolute value of the negative eigenvalue. The absolute values of all other negative eigenvalues constitute Lagrange multipliers of local minima. Granik further proposes to use the so-obtained (global and local) solutions as start points for a local optimization subject to the original box constraint. The box constraint allows transmission values ranging from -1 to 1 , which is suited, for example, for attenuated PSMs or CPL. In the same paper, a solution to the slightly more involved situation of a binary mask-suited constraint is shown.

As a third alternative for the image-fidelity-based mask optimization problem, Granik devises a direct non-linear problem formulation. Similarly to his source optimization approach, he reformulates (7.23) employing introducing a Sobolev norm for smoothing (see (7.10)):

$$\epsilon_w(\mathbf{O}) = (|w(\mathbf{I} - \mathbf{T})|^2 + \alpha_1 |L_1 \mathbf{O}|^2 + \alpha_2 |L_2 \mathbf{O}|^2 + \alpha_3 |\mathbf{O} - \mathbf{O}_0|^2)^{1/2}. \quad (7.22)$$

The operators $L_i (i = 1, 2)$ yield the respective derivatives whose squared Euclidean norm is weighted with the parameters α_i . In addition, the last term can be considered a regulation expression in that it can be used to incorporate intended mask properties. Specifically, it measures the difference between a design intent (\mathbf{O}_0) and the evaluated solution candidate. As before, a weight function w can be employed to emphasize or ignore regions in the deviation evaluation. In contrast to the previous QP formulation, the image difference is expressed in terms of a Euclidean norm, providing a smoother function compared with the Manhattan norm, because it turns the first term in (7.22) into a polynomial. The application of the Sobolev-norm scheme further improves the smoothness of the error function and hence leads to a reduction of discontinuities, especially in the derivatives of the error function. Granik states that despite an improved image fidelity, the so-posed figure of merit still exhibits contour errors. He thus extended the error function by a contour fidelity term (cf. Sayegh, 1982):

$$z_{\theta}(I, C) := \oint_C (I - \theta)^{2n} dz, \quad (7.23)$$

⁸Also see Nocedal and Wright (2006, pp. 70) for a discussion.

where I is the image, C is the intended contour and θ is a predefined threshold. The exponent n can be interpreted as a relaxation attribute for which Sayegh suggests values of 1 or 2. This leads to the following final optimization problem:

$$\begin{aligned} \arg \min_{\mathbf{O}} \quad & \epsilon_w(\mathbf{O}) = (|w(\mathbf{I} - \mathbf{T})|^2 + \alpha_1 |L_1 \mathbf{O}|^2 + \alpha_2 |L_2 \mathbf{O}|^2 + \alpha_3 |\mathbf{O} - \mathbf{O}_0|^2 + z_\theta(I, C))^{1/2} \\ \text{s.t.} \quad & |\mathbf{O}|_\infty - 1 \leq 0, \end{aligned} \quad (7.24)$$

Granik demonstrates the application of two different optimization methods to this problem: (1) a direct local variation method and (2) a gradient descent approach. The local variation method is similar to the iterative pixel flipping approach by Saleh and Sayegh (1981); Sayegh (1982). As schematically demonstrated in Algorithm 7.2, in each iteration the impact of each individual mask pixel is evaluated. This is done by determining the image error resulted by the original pixel value, and by an incremented and a decremented variation. The value leading to the smallest error is fixed. Pixels leading to a constraint violation are pulled back onto the boundary. After all pixels have been visited, both the increment and the decrement are contracted for the next iteration step. The procedure is repeated until a predefined termination criterion is met. The visitation order of the pixels is critical since changes in one pixel also influence the assessment outcome of neighboring ones. Without any contraction of the modification increments and decrements, the algorithm does in general not converge. Moreover, this procedure can be said to be “brute force” as it neither implicitly nor explicitly considers any information gathered in previous iteration steps. An alternative to this approach was more recently proposed by Zhang et al. (2008), aiming at reducing the number of evaluations by attempting to relate the impact of individual pixels to a global cost function.

Algorithm 7.2: Local variation optimization scheme.

```

Data:  $\mathbf{O}$ ; // initial mask
Data:  $\Delta_0 \leftarrow$ ; // initial mask pixel increment
1 while not termination criterion met do
2   foreach  $\mathbf{O}[l_x, l_y]$  do // for each mask element
3     for  $i = 0, 1, 2$  do // duplicate mask
4        $\mathbf{O}'_i \leftarrow \mathbf{O}$ ;
5        $\mathbf{O}'_1[l_x, l_y] \leftarrow \mathbf{O}[l_x, l_y] + \Delta^{(k)}$ ;
6        $\mathbf{O}'_2[l_x, l_y] \leftarrow \mathbf{O}[l_x, l_y] - \Delta^{(k)}$ ;
7       if  $|\mathbf{O}'_i[l_x, l_y]| > 1$  then
8         pull back;
9        $\mathbf{O} \leftarrow \arg \min_{\mathbf{O}'_i} \{\epsilon_w(\mathbf{O}'_i) \mid i = 0, \dots, 2\}$ ;
10       $\Delta^{(k+1)} \leftarrow \alpha \Delta^{(k)}$ ; // contraction ( $\alpha < 1$ )
11       $k \leftarrow k + 1$ ;

```

As an alternative, Granik proposes to use a gradient descent method, which if the merit function is sufficiently smooth and gradients can be obtained in a straightforward manner, i.e., analytically or efficiently numerically, and may lead to a much improved convergence behavior. Because of the local nature of the locality of quasi-Newton approaches, starting points have to be carefully chosen. According to Granik, an analytic solution of the quadratic problem, as demonstrated above, presents a viable initial point. Granik proposes the application of the *reduced gradient method* by Frank and Wolfe (1956) to solve the problem. This method, although not suited as an efficient general-purpose approach for the class of non-linear programs, exhibits certain advantages for large-scale optimization task, especially when a fast convergence to an approximate solution is intended. First order gradient approaches strongly depend on the efficient computation of the first derivatives. Granik derives an analytical approximation

as follows: Rewriting the first term of (7.25)—the image deviation measure—as the (weighted) inner product, we have that

$$\begin{aligned}\varepsilon_w^2(\mathbf{O} + \delta\mathbf{O}) &= \langle \mathbf{I} - \mathbf{T}, \mathbf{I} - \mathbf{T} \rangle \\ &\approx \varepsilon_w^2(\mathbf{O}) + 2\langle \mathbf{I} - \mathbf{T}, \delta\mathbf{I} \rangle,\end{aligned}\tag{7.25}$$

where the image derivative can be stated as follows:

$$\delta\mathbf{I} := \mathbf{I}_{\mathbf{O}+\delta\mathbf{O}} - \mathbf{I}_{\mathbf{O}},\tag{7.26}$$

which, when using the SOCS formulation becomes,

$$\begin{aligned}\delta\mathbf{I} &= \sum_m \lambda_m (\mathbf{x}_m * (\mathbf{O} + \delta\mathbf{O})) (\mathbf{x}_m * (\mathbf{O} + \delta\mathbf{O}))^* \\ &\quad - \sum_m \lambda_m (\mathbf{x}_m * \mathbf{O}) (\mathbf{x}_m * \mathbf{O})^* \\ &\approx \sum_m \lambda_m (\mathbf{A}_m (\mathbf{x}_m * \delta\mathbf{O}) + \mathbf{A}_m^* (\mathbf{x}_m * \delta\mathbf{O}))\end{aligned}\tag{7.27}$$

where \mathbf{A}_m can be regarded as the field amplitude under the m th kernel; formally,

$$\mathbf{A}_m := \mathbf{x}_m * \mathbf{O}.\tag{7.28}$$

Substituting the approximation of $\delta\mathbf{I}$ into the inner product $\langle \mathbf{I} - \mathbf{T}, \delta\mathbf{I} \rangle$, results:⁹

$$\begin{aligned}\langle \mathbf{I} - \mathbf{T}, \delta\mathbf{I} \rangle &\approx 2 \sum_m \lambda_m \Re \langle \mathbf{A}_m^* (\mathbf{I} - \mathbf{T}), \mathbf{x}_m * \delta\mathbf{O} \rangle \\ &= 2 \sum_m \lambda_m \Re \langle \delta\mathbf{O} * \mathbf{x}_m, \mathbf{A}_m^* (\mathbf{I} - \mathbf{T}) \rangle \\ &= 2 \sum_m \lambda_m \Re \langle \delta\mathbf{O}, \mathbf{x}_m^* \star w \mathbf{A}_m^* (\mathbf{I} - \mathbf{T}) \rangle,\end{aligned}\tag{7.30}$$

Thus with

$$\nabla \varepsilon_w^2 \approx \varepsilon_w^2(\mathbf{O} + \delta\mathbf{O}) - \varepsilon_w^2(\mathbf{O}) \approx 2\langle \mathbf{I} - \mathbf{T}, \delta\mathbf{I} \rangle_w\tag{7.31}$$

we have that

$$\nabla \varepsilon_w^2 \approx 4 \sum_m \lambda_m \Re(\mathbf{x}_m^* \star w \mathbf{A}_m^* (\mathbf{I} - \mathbf{T})).\tag{7.32}$$

According to Granik, this semi-analytical determination of the gradient leads to a significant performance increase as it reduces the dependence from the number of pixels in the image from $O(n^2)$ to $O(n \log n)$.

Another speed-up both for the local variation and the gradient descent approach can be achieved by exploiting the multi-linearity of the SOCS formulation and of convolution. In the spatial domain, the SOCS formula can be written as

$$\mathbf{I} = \sum_m \lambda_m \mathbf{A}_m \mathbf{A}_m^*.\tag{7.33}$$

Given a mask $\mathbf{O}^{(k)}$ for which the field amplitude under kernel m is known ($\mathbf{A}_m^{(k)}$). If the mask is only slightly changed—that is, the new mask ($\mathbf{O}^{(k)}$) exhibits only a small pixel difference from the

⁹Using that

$$\langle f * g, h \rangle_w \equiv \langle f^*, g \star wh \rangle,\tag{7.29}$$

where f, g and h are complex-valued functions and w is a weight factor for the inner product.

k th mask—the distributivity of convolution allows for a computation-time-efficient update scheme as follows:

$$\begin{aligned}
 \mathbf{A}_m^{(k+1)} &= \mathbf{x}_m * \mathbf{O}^{(k+1)} \\
 &= \mathbf{x}_m * (\mathbf{O}^{(k)} + \mathbf{O}^{(k+1)} - \mathbf{O}^{(k)}) \\
 &= \mathbf{x}_m * \mathbf{O}^{(k)} + \mathbf{x}_m * (\mathbf{O}^{(k+1)} - \mathbf{O}^{(k)}) \\
 &= \mathbf{A}_m^{(k)} + \mathbf{x}_m * (\mathbf{O}^{(k+1)} - \mathbf{O}^{(k)}).
 \end{aligned}$$

Granik demonstrates the applicability of the former scheme to the Frank-Wolfe method as follows: In each step of the method a linear program (LP) model of the original problem is solved, and the solution yields the search direction, for which the step length is computed through line search (also see (5.19)). Let \mathbf{d} be a solution to the LP, leading to the following line search problem:

$$\mathbf{O}^{(k+1)} := \mathbf{O}^{(k)} + \alpha \mathbf{d}. \quad (7.34)$$

Let, as before, $\mathbf{A}_m^{(k)}$ be the field amplitude due to the mask of the current iteration step. In addition, let $\mathbf{A}_m'^{(k+1)}$ be the amplitude under the LP solution. The amplitude due to the step length solution α can hence be obtained by the following linearization:

$$\begin{aligned}
 \mathbf{A}_m^{(k+1)} &= \mathbf{x}_m * (\mathbf{O}^{(k)} + \alpha \mathbf{d}) \\
 &= \mathbf{x}_m * \mathbf{O}^{(k)} + \mathbf{x}_m * \alpha \mathbf{d} \\
 &= (\mathbf{x}_m * \mathbf{O}^{(k)}) + \alpha (\mathbf{x}_m * \mathbf{d}) \\
 &= \mathbf{A}_m^{(k)} + \alpha \mathbf{A}_m'^{(k+1)}.
 \end{aligned} \quad (7.35)$$

Granik demonstrated the feasibility of his approach for the optimization of different feature set-ups including isolated contactholes, the contacthole layout example by Socha et al. (2004), and critical layers of random logic and SRAM cells. Significant improvements in terms of image and contour fidelity for both binary masks and PSMs. Because of the pixel representation, the proposed approach does not only provide classical OPC solution such as line-width variations or serifs but also “adds” SRAFs to the layout. However, the manufacturability criterion that can be incorporated into the merit function are limited. For example, later a restriction of scatter bars to rectilinear or 45-degree structures was implemented as a post-optimization regularization routine (Granik et al., 2006). Differences in terms of performance and mask write times were experimentally verified (Hendrickx et al., 2008). Recently, the approach was extended to high transmission AttPSMs (Sakajiri et al., 2009).

7.1.5 Local Mask Pattern Optimization

Another inverse lithography approach to source/mask optimization was proposed by Poonawala and Milanfar (2006, 2007c). The authors start with an abstract view on the lithographic process, in which the lithographic process is regarded as an operator T transforming an input mask pattern m into an intensity distribution or pattern representation z :

$$z(x, y) := T\{m(x, y)\}. \quad (7.36)$$

Consequently, inverse lithography is regarded as an inverse process, that, given a desired output, yields the corresponding mask. Thus the task is to find an inverse operator T^{-1} such that

$$\hat{m}(x, y) = T^{-1}\{z^*(x, y)\}. \quad (7.37)$$

The so-obtained ill-posed problem is to be stated as an optimization problem minimizing the difference between the target and the actual pattern:

$$\hat{m}(x, y) = \arg \min_{m(x, y)} d(z^*(x, y), T\{m(x, y)\}), \quad (7.38)$$

where d is the distance between a target pattern and an intensity output pattern achieved by a given input.

The transfer function T is defined as a continuous function, allowing for a gradient-based search approach. Both coherent and incoherent imaging systems can be taken into account, an extension to the partially coherent case—for example, using a (truncated) SOCS approach—is also indicated. A regularization framework is introduced in order to account for mask manufacturability criteria. Additionally, the aerial image performance in terms of contrast and side-lobes is evaluated and included in the merit function as penalty method. The transfer function is split into two parts: (1) the aerial image formation and (2) the resist processes. In their approach, the authors assume a simplified threshold-based resist model. In order for the proposed continuous function optimization approach to work, the discrete threshold model has to be transferred into a continuous one. Thus, the Heaviside operator is approximated by a sigmoid function, which is defined as follows:

$$\text{sig}(x) := \frac{1}{1 + \exp(-a(x - \theta))}. \quad (7.39)$$

Here, θ specifies the x value of the saddle point. This value corresponds to the given threshold intensity value of the resist. Above this threshold value, the positive photoresist is assumed to be completely removed. The slope of the sigmoid can be influence by the parameter a . Figure 7.3 demonstrates the impact of a .




Figure 7.3: Sigmoid function as defined in (7.39). The threshold value (θ) is fixed to 0.5. Two different slope values (parameter a) are illustrated.

The flow of the combined imaging and resist transfer procedure is depicted in Figure 7.4.




Figure 7.4: Flow of imaging and resist forward process.

In their work, the authors focus on mask optimization for fully coherent and fully incoherent imaging systems. Only binary and strong phase-shifting masks with phase values of 0 and π are regarded. Additionally, a pixelation of both input and output patterns is employed. Each pixel of the mask (m_j)

can hence be expressed as the real part of the complex transmission only. The minimization problem can so be reduced to:

$$\min \sum_j^N (z_k^* - z_k)^2, \quad (7.40)$$

where N is the number of pixels in both x and y direction. The electric field amplitude in the coherent imaging case can be described as convolution of the field of the mask U and a amplitude point spread function h :

$$I = |E|^2 = |U * h|^2, \quad (7.41)$$

where $x \mapsto |x|^2$ denotes the element-wise absolute square. The amplitude point spread function can be estimated by the Fourier transform of a circular aperture, which corresponds to J_1 , the first kind, first order Bessel function. Applying the resist threshold model to the intensity (square of electric field) and discretizing the convolution yields

$$z_k := \text{sig} \left[\left(\sum_j^N h_{kj} U_j \right)^2 \right] \quad (7.42)$$

for each pixel z_k .

In the incoherent case, the field amplitude is given by:

$$I = |E|^2 = |U|^2 * |h|^2, \quad (7.43)$$

where h again refers to the amplitude point spread function. Equivalently, $\tilde{h} := |h|^2$ is the point spread function. Thus,

$$z_k := \text{sig} \left[\left(\sum_j^N |h_{kj} U_j|^2 \right) \right], \quad (7.44)$$

for each pixel z_k . For binary masks, $|U_j|^2 = U_j$, thus (7.44), can be stated as follows:

$$z_k := \text{sig} \left[\left(\sum_j^N \tilde{h}_{kj} U_j \right) \right], \quad (7.45)$$

In order to apply a steepest-descent search algorithm to the minimization problem (7.40), the following gradient has to be computed:

$$\frac{\partial}{\partial U_j} \sum_j^N (z_k^* - z_k)^2, \quad (7.46)$$

which in the coherent case is equivalent to:

$$\begin{aligned}
& \frac{\partial}{\partial U_j} \sum_j^N \left(z_k^* - \text{sig} \left[\left(\sum_j^N h_{kj} U_j \right)^2 \right] \right)^2 \\
&= -4a \sum_{k=1}^N h_{kj} \left(\sum_{i=1}^N h_{ki} U_i \right) \exp \left(-a \left(\sum_{i=1}^N h_{ki} U_i \right)^2 + a\theta \right) \\
& \quad \left(z_k^* - \frac{1}{1 + \exp \left(-a \left(\sum_{i=1}^N h_{ki} U_i \right)^2 + a\theta \right)} \right) \\
& \quad \left(\frac{1}{1 + \exp \left(-a \left(\sum_{i=1}^N h_{ki} U_i \right)^2 + a\theta \right)} \right)^2 \\
&= -4a \sum_{k=1}^N h_{kj} \left(\sum_{i=1}^N h_{ki} U_i \right) \left(\frac{1}{\text{sig} \left[\left(\sum_{i=1}^N h_{ki} U_i \right)^2 \right]} - 1 \right) \\
& \quad \left(z_k^* - \text{sig} \left[\left(\sum_{i=1}^N h_{ki} U_i \right)^2 \right] \right) \\
& \quad \left(\text{sig} \left[\left(\sum_{i=1}^N h_{ki} U_i \right)^2 \right] \right)^2
\end{aligned} \tag{7.47}$$

Let $F = \sum_j^N (z_k^* - z_k)^2$:

$$\begin{aligned}
\frac{\partial F}{\partial U_j} &= -4a \sum_{k=1}^N h_{kj} \left(\sum_{i=1}^N h_{ki} U_i \right) \left(\frac{1}{z_k} - 1 \right) (z_k^* - z_k) (z_k)^2 \\
&= -4a \sum_{k=1}^N h_{kj} \left(\sum_{i=1}^N h_{ki} U_i \right) z_k (1 - z_k) (z_k^* - z_k)
\end{aligned} \tag{7.48}$$

The gradient can then be written as:

$$\nabla F(U) = -4a(H^T[(z^* - z) \circ z \circ (1 - z)](HU)), \tag{7.49}$$

where \circ is the entry-wise or Hadamard or Schur product. Similarly, in the incoherent case, the following expression can be obtained:

$$\nabla F(U) = -4a(H^T[(z^* - z) \circ z \circ (1 - z)]) \tag{7.50}$$

When aiming at binary masks, U_j is of course to be restricted values of either 0 or 1. Such a constraint, however, yields a discrete optimization problem. Since the authors were determined to use their continuous approach also for the binary mask, they introduce a box constraint:

$$\forall j \in [1..N])(0 \leq U_j \leq 1,$$

with an according mapping ($\mathbb{R} \rightarrow \mathbb{N}$). Additionally, to employ an unconstrained optimizer, the constrained is translated into a function mapping \mathbb{R} to the $[0, 1]$ range, for which the function

$$f : x \mapsto \frac{1}{2}(1 + \cos(x))$$

is used. The rationale for using this function is that due to its periodic nature, pixel values are flipped more frequently, therefore allowing for a more vivid search and a lower probability of a premature convergence. Being a strictly local approach, the devised method is only capable of yielding local optima and strongly depends on the selection of an adequate initial solution. A straightforward approach is, for example, the use of the target pattern. Albeit without the incorporation of additional criteria, such a choice can be expected to bias the search toward solutions close to the original design intent. An iteration step of the steepest descent algorithm is given by

$$\mathbf{x}^{n+1} := \mathbf{x}^n - s \nabla F(\mathbf{x}^n), \quad (7.51)$$

where s is the step size. The authors demonstrated the potential of this approach. In order to incorporate further constraints, for example, accounting for mask manufacturability, the authors developed penalty method approach, which they term regularization. Specifically, the authors introduced three penalty terms: (1) A *discretization penalty*, which introduces a bias favoring mask transmission values close to 0 or 1. (2) A *complexity penalty*, rewarding smooth mask, that is, those patterns that consist of larger regions with uniform transmission values. (3) And an *aerial image quality penalty*, favoring patterns with a smaller difference (for example in terms of an Euclidean norm distance measure) to the binary target pattern representations. This latter term is similar to the image fidelity term proposed by Granik (see Section 7.1.3) and is to some extent suited to ensure that solutions exhibit an improved local contrast. As pointed out by Granik (2006),¹⁰ the use of a Euclidean norm bears the risk of yielding solutions with both regions closely on target and those that are dramatically different from the intended image, a situation that could be circumvented by the use of the maximum norm. Poonawala and Milanfar (2007c) acknowledge this fact by stating that their image fidelity criterion should in general not be used without the main target, the contour fidelity.

As argued before, due to the strong locality of the proposed approach and the strong bias presented through the contour difference objective, solutions tend to be close to the nominal (non-optimized) mask layout. To attenuate this shortcoming Poonawala and Milanfar (2007a) devised a two-stage optimization strategy in which first only the image fidelity is maximized, thus removing the strong bias presented by the contour criterion. As a consequence the solutions can be expected to show a less pronounced resemblance with the target layout—in other words, show a higher level of innovation. The results of the former step are used as initial solutions for a subsequent optimization, accounting for the remaining criteria.

Albeit limited to only fully coherent and incoherent cases, the authors clearly demonstrated the feasibility and potential of their method, presenting well-performing solutions for different mask and patterning technologies (e.g., Poonawala and Milanfar, 2007b). As common for penalty methods, the selection of weight parameters strongly impacts the resulted solutions. The authors presented a number of comparisons between different weight settings vividly demonstrating this fact.

The partially coherent illumination case indicated by the original authors was later further developed by Ma and Arce (2008b). In their extension, aimed at CPL, only the SOCS kernel corresponding to the largest eigenvalue is considered,¹¹ allowing for a similarly concise derivation of gradients as in the fully coherent case. Using the same framework, the authors also demonstrated the optimization of binary masks (Ma and Arce, 2008a).

Another extension developed by Ma and Arce (2007) further extended the framework by introducing an alternative regularization term based on a Haar wavelet transform of the mask and by which they penalize masks with high frequency patterns (i.e., patterns with a high level of detail). The penalty can be applied either globally or locally, such that regions can be discriminated according to the detail level that they may exhibit.

A closely related approach that rigorously accounts for partial coherence was very recently devised by Jia and Lam (2011). There, process variations are probabilistically taken into account, rendering it less prone to yielding solutions non-feasible under realistic process conditions. Another interesting path studied by the same authors is the application of an incremental learning approach to a through-focus

¹⁰And a similar argument can also be found in Rosenbluth et al. (2002).

¹¹Following similar considerations as, for example, Socha et al. (2004).

mask optimization problem (Jia and Lam, 2010). There, the set of images at different focus positions is regarded as the training set. One probabilistically selected defocused image is considered at one iteration step.¹²

7.1.6 Level-set-function-based Approach

Another innovative ILT technique was developed by a group of authors (see, for example, Liu et al., 2005; Abrams and Pang, 2006; Pang et al., 2007). Like in the previously discussed approach (see Section 7.1.5), an abstract inverse problem formulation is derived, assuming an operation T transforming a mask function into a wafer pattern (7.36). Straightforwardly, the inverse is then defined as a transformation from a target pattern to a mask pattern that produces the target (7.37). As an example optimization target, replacing the ill-posed inverse problem, Abrams and Pang (2006) propose a total variation distance metric:

$$H(\phi) = \iint_{x,y} |T(\phi) - z^*| dx dy, \quad (7.52)$$

where ϕ is the mask function and T and z^* are the pattern forming operation and the target pattern, respectively. This formulation does not account for process conditions such as focus variations or image fidelity considerations. Thus the operator T transforms one mask pattern into exactly one wafer pattern. Many of process window related aspects can be modeled into this merit function in a straight-forward manner. For example, different focus positions could be taken into account by forming a weighted sum over the differences yielded by the different image results for different focus positions. Other criteria require a sophisticated formulation. For example, the mask manufacturability may be considered both as a figure merit—simpler mask receive a lower error value—or as a constraint, omitting non-producible mask functions from the considered set. The main advantage of a so-defined merit function is that it can be regarded as a Hamiltonian mapping a two-dimensional function (out of some appropriate Hilbert space) to a real value, the error. This property renders the problem amenable to a level-set formulation (Osher and Sethian, 1988; Sethian and Adalsteinsson, 1997),¹³ in which the mask function ϕ is defined as the level-set function, greater than zero inside, less than zero outside and zero on the contour of the mask. In other words, the contour of the mask is given as the set of coordinates for which x, y is zero. The optimization task is to find a level-set function that minimizes the error due to H , which is equivalent to finding ϕ such that the functional derivative of H with respect to ϕ becomes zero, numerically solvable, for example, using gradient-based optimization techniques.¹⁴

The level-set formulation presents a number of expedient characteristics. The functional representation of the mask, for instance, allows for a reduced of the parameters compared to traditional OPC or other SMO approaches. Moreover, due to the implicit retrieval of contours, a number of topography-related constraints can be accounted for in a straightforward fashion. Consequently, SRAF placement or, more generally, the decomposition of an individual feature into several sub-features can be considered a “natural” property of the level-set approach. This is demonstrated in the simplistic example of Figure 7.5. There, following a common practice, the level-set function of a fictive mask is defined as the signed distance of points to the closest edge. A simple translation in the z -direction leads to a significant topological modification of the pattern. The example also demonstrates that the volatile geometries require an additional effort to ensure the producibility of the attained masks, best seen when noting the small squares of the “S” in Figure 7.5(e).

The major drawback of this method is its limited flexibility in terms of joint source optimization. This has been alleviated, however, by introducing an additional level-set function for the source. There, similarly as for the mask optimization task, contours representing source regions are generated through the evolution of a level-set function (e.g., Pang et al., 2010). Like for the mask, regularizing constraints

¹²Also see Section 5.6.1.

¹³Level-set and related methods are commonly used in front-tracking or similar applications such as etching or photoresist development, see, for example, Section 3.4.

¹⁴A concise mathematical derivation of a similar approach, albeit restricted to fully coherent sources, was recently presented by Shen et al. (2009).




Figure 7.5: Signed distance function representation of multiple feature contours under z -translation: (a) -15, (b) -5, (c) 0, (d) 5 and (e) 10, which result in different topologies.

can be included to account for producibility or robustness requirements of the source. Additional objectives such as process-window-related criteria (ibid.) or MEEF (Pang et al., 2008b) and aberration and polarization-induced effects (Dam et al., 2010) have also been shown to be integrable into the approach. Multiple results demonstrate the feasibility and the benefits of the presented approach using mask-only optimization, for example for DRAM structures (Chu et al., 2006) on the 65 nm technology node (Hung et al., 2006), and for the 45 nm and 32 nm nodes (Pang et al., 2007, 2008a).

7.1.7 Pixelated Chromeless Phase Lithography Mask Optimization

Another comprehensive mask optimization framework was proposed by a group of authors (Singh et al., 2008; Davids and Bollepalli, 2008), aiming mainly at CPL¹⁵ (Borodovsky et al., 2008; Schenker et al., 2008) but being principally extensible to any other mask technology (Singh et al., 2008). One of the key techniques of their optimization approach is the use of an indirect representation of the mask that the authors refer to as Tau-map. Similarly to the previous discussed wavefront engineering approaches (see Section 7.1.1, Section 7.1.2 and Section 7.1.3), the basic idea is to optimize the Tau-map in order to obtain an ideal imaging performance. Summarizing the paper by (Davids and Bollepalli, 2008), the main ideas, (1) an illumination-independent forward imaging step as a function of the Tau-map, (2) the formulations for the according Tau-map optimization problem, and (3) the reverse mask construction algorithm, are briefly reviewed in the following.

As indicated before, the mask transmission and phase values of the mask are given on the discrete Tau-map, whose resolution is determined by the optical system, that is, by the NA of the projector lens, the employed wavelength and the filling factor of the illuminator. Of course, the sampling rate has to be greater than or equal to the Nyquist number. The Tau-map can hence be perceived as a mask spectrum with a resolution perfectly matching the propagation capability of the optical system. The point-spread function of the projector serves as an interpolation function that reconstructs the mask spectrum within its support, i.e., within the NA of the pupil.

Figure 7.8 demonstrates the Tau-map and the image computation scheme: An individual image point \mathbf{r}' (mask scale) is evaluated in the grid coordinate system. The imaging behavior of the point-spread function (U_1) on an individual Tau-map entry can be expressed as a weighting equation:

$$\mathfrak{D}(\mathbf{r}'; n, m) := \tau[n, m]U_1(k\text{NA}|\mathbf{r}' - \mathbf{x}_{n,m}|), \quad (7.53)$$

where k is the wavenumber and $\mathbf{x}_{n,m}$ is the actual coordinate corresponding to the Tau-map entry n, m . Here, the point-spread function can be said to be both a weighting and an interpolation function. Under partially coherent illumination, the mutual coherence of the imaged Tau-map entries has to be taken into account (cf. Section 3.2). This can be achieved by a properly stacked mutual intensity matrix (\mathbf{J}) defined on the same grid as the Tau-map. We then have the following equation for the image intensity:

$$I(\mathbf{r}') = I_0 \sum_n \sum_m \sum_{n'} \sum_{m'} \mathfrak{D}^*(\mathbf{r}'; n, m) \mathbf{J}[n, m; n', m'] \mathfrak{D}(\mathbf{r}'; n', m') \quad (7.54)$$

The intensity computation can be further simplified by using a translated coordinate system centered at the coordinate $\mathbf{x}_{a,b}$ that is closest to the evaluation point. Moreover, because of the band-limiting property of the pupil, only a portion of the grid has to be taken into account. Let $\pm l_x, \pm l_y$ be the indices corresponding to the NA-determined limits. If we set $\boldsymbol{\rho} := \mathbf{r}' - \mathbf{x}_{a,b}$, we obtain:

$$I(\mathbf{r}') = I_{a,b}(\boldsymbol{\rho}) = I_0 \sum_{\nu=-l_x}^{l_x} \sum_{\mu=-l_y}^{l_y} \sum_{\nu'=-l_x}^{l_x} \sum_{\mu'=-l_y}^{l_y} \mathfrak{D}^*(\boldsymbol{\rho}; \nu + a, \mu + b) \mathbf{J}[\nu, \mu; \nu', \mu'] \mathfrak{D}(\boldsymbol{\rho}; \nu' + a, \mu' + b). \quad (7.55)$$

This simplification is valid because, as stated before (Section 3.2), the mutual intensity depends on the difference between two points, not on its absolute values. It is hence invariant to a translation. Moreover, because the mutual intensity matrix is real and symmetric (for symmetric illumination sources), and

¹⁵See Section 2.6.3.




Figure 7.8: (a) Tau-map: sampled mask spectrum representation with a grid size and discretization determined by the wavelength, the illuminator and the NA of the projector. In addition, the pupil implies a band-pass filter, so that only the gray-shaded region, centered at the point where the image is evaluated has to be taken into account. For image points with no direct correspondence in the Tau-map, the point-spread function serves as an interpolation function. (b) Pupil frequency and illumination domain (normalized by the wavenumber $[k]$): the gray-shaded circle denotes the NA of the pupil, which corresponds to the support in the Tau-map.

hence hermitian, an eigen-decomposition using an orthogonal similarity transform can be carried out such that

$$\mathbf{\Lambda} = \mathbf{V}^T \mathbf{J} \mathbf{V}, \quad (7.56)$$

where $\mathbf{\Lambda}$ is a diagonal matrix containing the eigenvalues and \mathbf{V} is an orthogonal matrix whose columns are the eigenvectors. This decomposition gives rise to a formulation that is closely related to the SOCS formulation of the Hopkins imaging approach (see Section 3.2.6). In this case, however, the coherent kernels can be regarded as the pseudo-electric fields due to the Tau-map filtered by the PSF that is weighted by the source (Davids and Bollepalli, 2008). Formally,

$$\mathfrak{E}_{a,b}(\boldsymbol{\rho}; \nu, \mu) := \sum_{\nu'=-l_x}^{l_x} \sum_{\mu'=-l_y}^{l_y} \tau[\nu' + a, \mu' + b] \mathbf{V}[\nu, \mu; \nu', \mu'] U_1(k \text{ NA} |\boldsymbol{\rho} - \mathbf{x}_{\nu', \mu'}|), \quad (7.57)$$

which corresponds to a discrete convolution of the Tau-map with the mutual-intensity-weighted PSF,¹⁶ and by accordingly substituting it into (7.55), we have:

$$I_{a,b}(\boldsymbol{\rho}) = I_0 \sum_{\nu=-l_x}^{l_x} \sum_{\mu=-l_y}^{l_y} \lambda_{\nu, \mu} |\mathfrak{E}_{a,b}(\boldsymbol{\rho}, \nu, \mu)|^2. \quad (7.58)$$

Based on this compact forward imaging formulation, the inverse step can be described as a search for a Tau-map that minimizes the difference between the target image and the image due to the Tau-map. This problem can be solved with a number of different optimization schemes. The authors schematically demonstrate the application of two approaches: (1) a first-order technique in which each iterate consists of an LP sub-problem and (2) a second-order weighted least squares formulation.

¹⁶This property becomes more obvious when replacing $\tau[\nu' + a, \mu' + b]$ with $\tau[\nu' - a, \mu' - b]$, which because of the symmetric sum limits just corresponds to a reordering.

For the first approach, termed linear variation techniques by the authors, Tau-map variations are related to the difference between the target (I^*) and the image under the map (I) as first order derivatives. That is, given

$$\Delta I := I^* - I,$$

ΔI has to be formulated as a function of the Tau-map. From (7.55), restricted to the actual sampling grid, that is omitting the interpolation vector $\boldsymbol{\rho}$, Davids and Bollepalli (2008) derive the following differential expression:

$$\begin{aligned} \Delta I_{a,b} = I_0 \sum_{\nu=-l_x}^{l_x} \sum_{\mu=-l_y}^{l_y} \sum_{\nu'=-l_x}^{l_x} \sum_{\mu'=-l_y}^{l_y} \Delta \mathfrak{D}^*(\nu + a, \mu + b) \mathbf{J}[\nu, \mu; \nu', \mu'] \mathfrak{D}(\nu' + a, \mu' + b) \\ + \mathfrak{D}^*(\nu + a, \mu + b) \mathbf{J}[\nu, \mu; \nu', \mu'] \Delta \mathfrak{D}(\nu' + a, \mu' + b) + R_{a,b}, \end{aligned} \quad (7.59)$$

where $R_{a,b}$ is a residual error. From (7.59), a linear approximation for the Tau-map derivatives, which are directly proportionate to the pseudo-electric field differentials, can be derived and applied in a line-search. As noted by Davids and Bollepalli (2008), the computation of the derivatives requires the inverse of the mutual intensity matrix \mathbf{J} , which can be obtained through a diagonalization similarly to (7.56), potentially involving a preconditioning to a singular eigenvalue matrix \mathbf{D} .

For the second approach, a weighted least squares regression, the authors propose the solution with a quasi-Newton method, where the sub-problem is obtained through a Taylor series expansion about the iterative candidate solutions of the Tau-map, consequently requiring the numerical determination of the Hessian of the functional defined by the residual error on the grid. As they point out, the Hessian is in general ill-conditioned, requiring according preconditioning (compare Section 5.4.1). As an alternative, the authors propose the use of smoothing techniques such as the introduction of a Sobolev norm (similarly, for example, to the approach taken by Granik (2004); see Section 7.1.3) or different search strategies such as direct search (e.g., Section 5.4.2).

As the final step, the reconstruction from the Tau-map obtained from either optimization procedure remains to be conducted. Like already mentioned in the discussion, for example, of the approach by Rosenbluth et al. (2002), this is again an ill-posed problem since numerous mask layouts may result in the same diffraction pattern. This is because of the band-limiting behavior of the lens, but it becomes even more critical when considering mask topography effects. The set of potential solutions is therefore best restricted using additional criteria such as mask manufacturing constraints, for example, by requiring a minimum pixel size that reflects the mask writer resolution capabilities. Unfortunately, by introducing restrictions not taken into account in the previous optimization stage, such an approach easily leads to sub-optimal solutions.

Albeit not explicitly shown, the authors claim that their approach allows to account for vector imaging, resist and EMF effects. However, such extensions are likely to break the presented concise representation and can be expected to be associated with a significant computational overhead. To alleviate the computational cost of rigorous EMF calculations, the use of a domain-decomposition technique (see Section 3.1.1) or approximation models such as the boundary layer method (Tirapu-Azpiroz and Burchard, 2003) have been proposed.

Of course, the Tau-map formulation allows to be used within virtually any including a global optimization framework. When combined with a local search, however, it strongly depends on the provision of an initial solution that lies within the basin of attraction of a suited local optimum. This strong dependence was also clearly demonstrated by Singh et al. (2008), in which different initial pixelations are discussed. Singh et al. (2008) also briefly discuss the problems involved when regarding source/mask optimization as an inverse lithography approach. He expresses this as a problem of generating an adequate target image that serves as comparison with the image obtained through an optimization step. To account for additional criteria such as contrast and image fidelity, the authors propose to include according parameters into the computation of the target image, indicating that a number of experiments had to be conducted in order to find suited such parameters. The feasibility of the approach was demonstrated up to full-chip scale through a number of other numerical experiments and silicon validations (ibid.).

7.2 Direct Problem Statement

In contrast to many of the related efforts to mutually optimize lithographic mask and illumination settings, using an inverse problem formulation, the approach proposed here can be considered a direct method, taking a holistic view on the lithography process.¹⁷ Instead of starting with an aerial image or the mask's nearfield and searching for (closely) matching mask and illumination set-ups, this technique (probabilistically) searches the entire space of all possible mask and illumination set-ups in order to find an “ideal” combination. Ideal in this case can be defined very flexibly, including virtually any criteria, for example, image performance related objectives or producibility constraints. The number and kind of criteria can be adjusted with very little effort. Furthermore, for the aerial (or bulk) image evaluation, no specially tailored simulation routines have to be used. Instead a normal lithography process simulation can be conducted. Thus also virtually any electro-magnetic field solver can be utilized, which will be a requirement for next generation lithography applications.

Figure 7.9 depicts the general flow of the proposed optimization approach. It also demonstrates, that this procedure can be easily integrated into the proposed optimizer, the genetic algorithm. Furthermore, the lithography simulation environment can be applied for the computation and evaluation of aerial (or bulk) images in a straightforward manner: (1) Mask and illumination settings are generated by the genetic algorithm. (2) Resulted solution candidates are evaluated using a simulation program and a post-processing analysis. (3) The obtained information (parameter settings and merit values) are used by the optimization routine to generate new solutions. These iteration steps are repeated until a satisfactory solution is found.

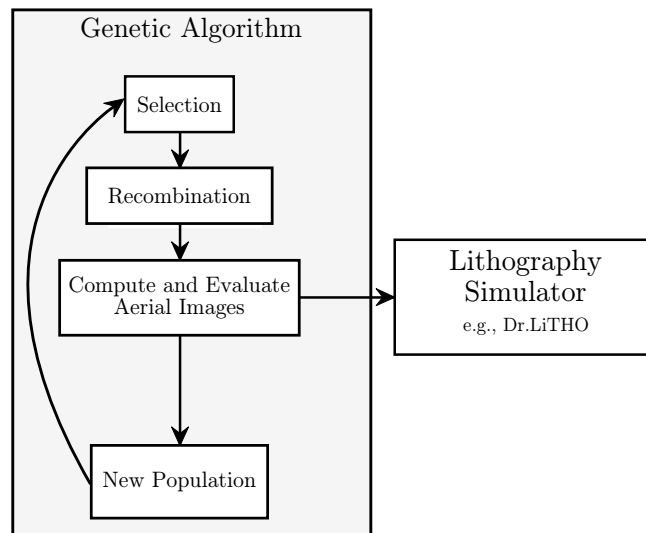


Figure 7.9: General optimization flow.

As indicated before, in genetic algorithms there is a dualism between the internal representation of solutions and the actual problem space: In this application, the problem specific representation (phenotype) of mask and illumination settings has to correspond with the required input for the lithography simulator. For the internal GA representation (the genotype) both mask and illumination parameters have to be represented as bit strings.

Another requirement for this optimization approach is the formulation of the merit (fitness) function. As in most other optimization problems, the pay-off of a solution is given by a number of (often incommensurable) different criteria. For the single-objective case, all objectives that determine the merit of a solution have to be modeled into a function that yields a single scalar value. This value serves as the

¹⁷Parts of this section have been previously published in (Fühner et al., 2007a).

fitness of a solution candidate for the genetic algorithm. In the case of a multi-objective optimization GA, sub-criteria can be evaluated individually, as the optimizer will treat them multi-dimensionally.

The parameter representation, the objective criteria, and the formulation of the merit function are presented in the remainder of this section.

7.2.1 Parameter Representation

Several different representation schemes for both the mask geometry and the illumination set-up have been investigated. In general, two categories of values can be distinguished: (1) continuous and (2) discrete.

Continuous values, such as the parameters for the continuous mask geometry representation and the illumination system are modeled as follows: The parameter's domain, that is, upper and lower bounds of the parameter (u and l) and the required number of decimal places (d), is specified in advance. The number of bits taken by this parameter in the chromosome can then be calculated by shifting the parameter's domain ($u - l$) by the required number of decimal places, and computing the required (up-rounded) exponent of radix two:

$$\ell := \lceil \log_2(u - l)10^d + 1 \rceil.$$

Any real number x can now be encoded by subtracting the offset (lower bound l) and dividing it by the range of the parameter ($u - l$); yielding numbers ranging from 0 to 1. The result is then scaled by the maximum number that can be represented by this parameter's part of the bit string ($2^\ell - 1$): $a_{10}(x) := \frac{x-l}{u-l} \cdot (2^\ell - 1)$. The binary representation ($a_2(x) := \text{bin}(a_{10}(x))$), finally, yields the parameter's bit encoding in the chromosome. Thus, decoding of the parameter is performed as follows: $x := \frac{\text{dec}(a_2)}{2^\ell - 1}(u - l) + l$.

Discrete values, such as transmission levels, are encoded bit-wisely.

Photomask

In the first examples, we have limited our view to the Kirchhoff mask approach (infinitely thin masks). Hence the mask geometry is defined in one or two dimensions. The one-dimensional representation is well suited for the determination of ideal conditions for line/space patterns, whereas for more complicated structures—such as contact holes—or an evaluation of line-end properties a two-dimensional representation is required. In both cases, the geometry can be represented continuously, that is, by a parametric geometry description, or discretely, in the form of a grid.

One-dimensional Masks The optimization problem for one-dimensional mask with line/space patterns is specified in two different ways. In the first case, the mask is represented by continuous variables which describe certain features on the mask. Figure 7.10a depicts a main feature with sub-resolution assists, which are symmetrically positioned on the left and on the right side of the main feature. Simulations are performed using different numbers of assist features.

Photomasks in optical lithography are commonly generated by mask writers (e-beam or laser), which have a limited resolution. This is directly taken into account for the second, the pixelated masks representation (demonstrated in Figure 7.10b). The mask is divided into pixels of size Δx that is determined by the resolution of the mask writer. Each pixel takes predefined complex transmission values t_i . In the binary mask (BIM) case, two values occur:

$$\begin{aligned} t_1 &: 0 + 0j \\ t_2 &: 1 + 0j. \end{aligned}$$

Alternating and attenuated phase shifting mask pixels may take one of the following three values:

$$\begin{aligned} t_1 &: 0 + 0j \\ t_2 &: 1 + 0j \\ t_3 &: -1 + 0j. \end{aligned}$$

Parameters for the discrete mask representation are converted in a straightforward manner. For binary masks, each grid point can only take two transmission values, thus requiring only a single bit. Depending on the technology, phase shifting masks require two (CPL) or three values (both AttPSM or AltPSM). While for the former case two bits are sufficient, the latter cannot be coded with an integer number of bits. There are multiple options to circumvent this problem, for example: (1) introducing tertiary alphabet, (2) disrespecting bit boundaries and aggregating multiple pixels into a number of integers, or (3) associating multiple bit configuration with a single (complex) transmission. For two reasons, we employed the latter. First, because it is compatible with the rest of our (binary) chromosome and second since an interweaving of pixel boundaries can be expected to introduce dispensable epistasis. Additionally, the non-unique assignment offers the possibility to introduce a bias toward a specific transmission, for example, to incorporate an implicit preference of chrome regions.




Figure 7.10: Mask geometry representations: Gray areas symbolize opaque chromium covered areas of the mask; continuous (a) and discrete (b) representations are illustrated for the one-dimensional case; w_m —width of main feature, w_{ai} —width of assist feature i , d_i —distance of assist feature i , Δx —grid resolution. For the two-dimensional rectangle representation (c) only the structures of the upper-right quadrant are defined; the remaining quadrants are mirrored.

Two-dimensional Masks Both the continuous and the discrete approach can be generalized to two dimensions. However, both representations have certain advantages over each other: The binary representation is very flexible and reflects the procedure of the mask production. On the other hand, it converges much slower than its counterpart, the continuous representation, which—due to the fixed number of assist structures—is associated with a significantly smaller search space, but which is hence also less flexible.

Therefore a third representation was developed as a compromise between the former approaches: The mask is described by a list of rectangles (Figure 7.10c). Each rectangle is characterized by four parameters (either origin, width and height or two opposite vertices). The number of rectangles may vary up to a user-provided maximum. The rectangle mask representation is flexible and capable to “invent” new layout-forms. Unfortunately, the varying number of rectangles (and important optimization parameters) limits the convergence of the implemented optimization procedure. Section 7.3.2 will present the application of the rectangle mask representation for the optimization of mask and source geometries for printing contact holes.

Illumination

Under Köhler illumination, the mask can be considered to be illuminated by plane waves with different directions (see Chapter 2). The Abbe approach employed in this work (see Section 3.2.5), accounts for that by computing the intensity distribution for each of these incident waves, which are then incoherently superposed to form the final image.

The traditional source geometries are revisited in Figure 7.11. In a first approach, the source is represented by a parametric definition of these illumination set-ups: The circular and annular illumination are characterized by one (σ) or two parameters (σ_{in} , σ_{out}), respectively. The multipole illumination is specified by the number of poles, by their orientation (θ), the offset (σ_{off}), and the radius (σ_r) of the poles. The implementation of other illumination types such as QUASAR or even user-defined geometries is straightforward.




Figure 7.11: Illumination models.

Since in this approach the principal illumination layout has to be pre-specified by the user, an alternative more flexible approach was developed, using a pixel representation of the source: As depicted in Figure 7.12, the source is sub-divided into tracks and sectors, each of which can have a different transmission value. Symmetry is exploited by restricting the source definition to only one quadrant, which is mirrored accordingly. Each track/sector can be regarded as one bit, which is either set (transmission = 1) or not (transmission = 0). The GA representation is thus straightforward: Figure 7.12(c), for example, shows a setup for bit string 000010000 as only the second sector in the second track has transmission 1, all other sections have 0 transmission.

In the course of the simulation studies, different Abbe-based aerial image simulators have been applied. For the first studies, a variant using a fixed mesh grid for the source was used into which the track/sector representation has to be expressed. Figure 7.12(d) shows the simulator's representation of the former example.

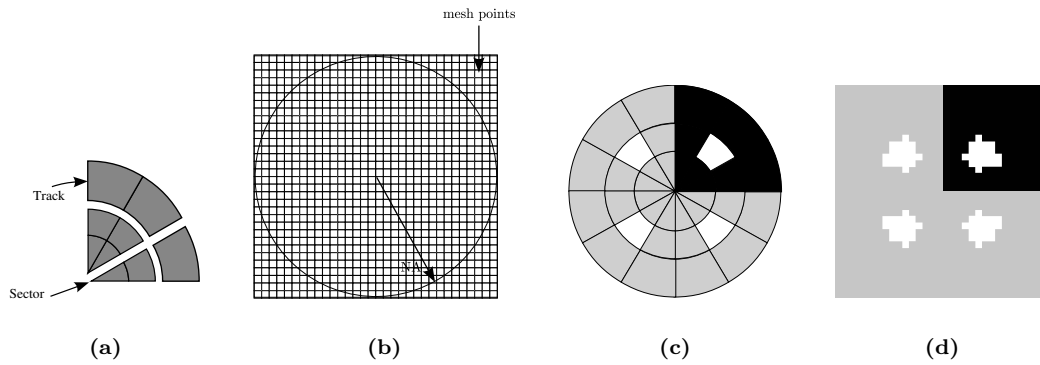


Figure 7.12: Illumination representation. (a) Tracks and Sector definition used by the genetic algorithm, (b) Pupil mesh point representation as used by lithography simulator: Matrix is specified that envelops the NA, (c) example set-up, (d) same set-up as mesh point representation

Figure 7.13 exemplifies the components of a bit string, given the track/sector illumination and rectangle representation of the mask.




Figure 7.13: Bit string representation of illumination parameters and rectangles.

7.2.2 Objectives and Constraints

It should be mentioned that the considerations in this work are solely aerial-image and bulk-image-based. That is, no resist effects have been taken into account. Instead, a constant threshold resist model is employed. There, an image computation is performed with a nominal line/space pattern, that is assumed to be on target without any OPC measures. In other words, mask and wafer linewidth are assumed to be identical. The intensity threshold for which this is the case is sought and kept constant for all other patterns.

As already illustrated through the review of related work, the merit of a solution is determined by a number of different criteria, related to the performance of a candidate solution but also accounting for manufacturability constraints. The discussion in Section 4.4 showed that constraint-handling in GAs is far from being straightforward such that replacement formulations have to be devised. As will be demonstrated, in this work, we have used a number of constraint handling techniques, including the penalty method, a multi-objective approach, a decoder and a repair strategy.

Another perspective on constraints is their impact on the solution. Hard constraints are those whose violations render a solution completely “useless.” In the case of soft constraints, a lower degree of violation enhances the utility of that solution; there may or may not be a turning point at which the solution is actually considered infeasible. These kind of constraints can alternatively be considered as regularization criteria, increasing the uniqueness of solutions of ill-posed problems. In this study, we employed criteria that can be regarded as hybrids incorporating both the pay-off and the degree of constraint violation. In other words, these figures of merit function as penalties, regularization criteria and pay-off values. For example, in order to minimize the risk of sidelobe printability, especially in the case of SRAFs, an according criterion can be formulated as follows: If the nominal threshold leads to sidelobe printing, the candidate solution is punished by a value proportionate to the difference between the nominal threshold and a safe intensity at which sidelobes do not print. If the nominal threshold results in no sidelobe, the solution is rewarded proportionate to the threshold-to-sidelobe distance. Hence since we are aiming at a straightforward and concise formulation as the merit function (single-objective or multi-objective), we will not draw a strict line between objectives and constraints but instead term all components of the merit criteria.

Broadly, the evaluation of an individual involves the following steps:

- (1) mask evaluation/repair,
- (2) image computation and
- (3) image evaluation.

The criteria are hence introduced in this order in the following.

Manufacturability

In order to achieve producible and inspectable masks, a number manufacturability criterion were introduced. The merit evaluation for this criterion strongly depends on the mask representation.

One-dimensional Masks Since the simulator employed for the first example applications represents mask elements by rectangles, the continuous and discrete representations require an extra step: The mask structures have to be converted to rectangles, which in both cases can be done in a straightforward manner. Since in the continuous case, parameters can only be varied such that they yield valid and producible masks, this representation does not require any additional geometry check. The binary representation, however, may generate cleft masks, i.e., masks that have a high number of adjacent regions with alternating transmissions. As such masks are hard to produce, higher merit values are assigned to masks exhibiting large connected areas with the same transmission (manufacturability criterion [MC]).

Two-dimensional Masks The manufacturability criterion is also applied in the case of the two dimensional rectangle representation. In this case, the validity of the rectangles has to be checked and the rectangles have to be modified such that they are processable by the lithography simulator.




Figure 7.14: Overlap healing. (a) initial condition: overlapping rectangles, (b) union—yields rectilinear polygon, (c) repaired (non-overlapping) rectangles

Again dealing with a precondition of one of the employed simulators, overlapping rectangles, potentially obtained as candidate solutions, have to be resolved. Thus the overlapping does not only have to be taken into account for the merit, but the overlap has to be resolved, that is, a combination of a repair and a penalty method is utilized. Two alternatives of dealing with overlapping structures have been implemented: (a) deletion of overlapping rectangles, (b) re-structuring. In the first alternative (a), for each rectangle, it is checked whether other rectangles intersect this rectangle. Each rectangle that intersects the checked rectangle is removed. The re-structuring alternative (b) is illustrated in Figure 7.14. First, the union of all overlapping rectangles is generated, resulting in a rectilinear polygon. This polygon is divided into rectangles afterward (for example, using the algorithm by Nahar and Sahni, 1988). The fitness is decreased by a penalty which is proportional to the overlap areas.

In a second step, rectangles that do not meet a predefined minimum size (Figure 7.15a) are punished: A value proportional to the difference of minimum size and actual size is subtracted from the merit; optionally, these rectangles are also deleted, thus again employing a hybrid of a repair and a penalty approach.

The decode function, which converts the bit string into rectangles, is not surjective, that is, there may be multiple definitions for the same mask set-up. For example, the rectangle defined by points (1, 4) and (4, 1) is equal to rectangle (1, 1), (4, 4). Also, rectangles may occur in a different order (in the bit string), still generating the same layout: ... R_1, R_2 ...and ... R_2, R_1 ... (where R_1 and R_2 denote rectangles) would lead to the same geometry. In order to improve the GA's convergence behavior, layouts with different bit string representations have been made distinguishable by punishing settings where the bit string representation does not reflect the geometric position. The example in Figure 7.15b would be punished, as the second rectangle in the bit string (denoted by the number 2) has a smaller distance (from the origin) than the first rectangle.

If rectangles lie outside of the specified mask area, they are punished proportionally to the area by which they exceed the mask array (this criterion is relevant only if rectangles are specified by origins and dimensions; Figure 7.15c). For manufacturability reasons, rectangles must respect a pre-specified minimum distance. Rectangles that under-run this distance are punished by a proportion of the violation (Figure 7.15d). Punished rectangles are deleted optionally (that is, the layout is repaired).

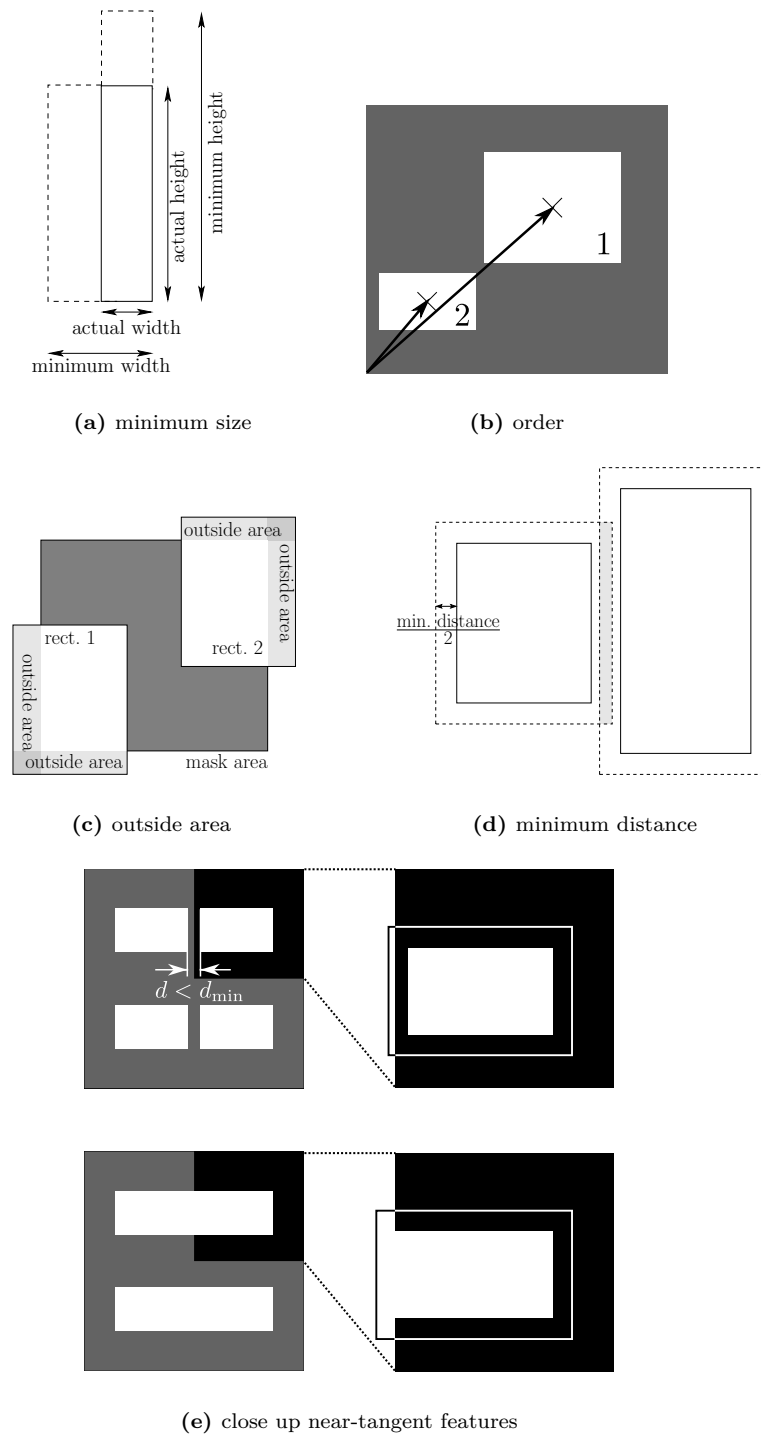


Figure 7.15: Rectangle representation: different criteria and repair steps are applied in the “geometry stage.” (a) Features below a minimum dimension are punished and expanded to the minimum size; (b) rectangles whose order within the candidate representation does not reflect the distance from the origin of the mask are punished; (c) rectangles outside of the mask area are repaired and punished; (d) rectangles whose distance to adjacent features falls below a minimum are punished and optionally deleted; (e) exploiting symmetry, the mask is defined only by one quadrant (e.g., upper right, shown magnified). The remaining quadrants are implicitly given as the mirrored representatives. Minimum distance violations at a symmetry axis are punished, and rectangles are rectified by pulling them onto the respective edge such that the symmetric instances form a single feature.

For symmetry reasons—the first 2-D application aims at optimizing masks for contact holes—the mask is defined by its upper right quadrant. The remaining quadrants are mirrored accordingly. Thus rectangles must have at least half the minimum distance from the left and bottom border of the mask area, as they would get too close to their mirrored instance otherwise. Rectangles violating this criterion are stretched (repaired) such that their offending edge then lies on the border (Figure 7.15e). Additionally, they are punished with a proportion of their distance violation.

Morphological mask evaluation An alternative method with a closer resemblance to the mask representation used in mask writing was devised by Seifert (2006b), borrowed from a morphological regularization technique used in image processing (see, for example, Haralick et al., 1987). Similarly to the mask manufacturability criteria above, the rationale behind this is to minimize the occurrence of (1) isolated structures that may be critical for the mask writer to achieve and that increase the writing time and of (2) structures with too low a distance in between, for which the writer may produce unintended bridges. In the first step, a binary representation of the masks is obtained. The pixelation is performed for a fixed grid, for example, with a resolution adapted to the mask writer capabilities. Let I_m be the binary map for mask m ; formally,

$$I_m : E \rightarrow T, \quad (7.60)$$

where E is the fixed mesh, i.e., $E \subset \mathbb{Z}^2$ and T is a set of transmission values. For simplicity, let T be $\{0, 1\}$, i.e., a binary mask. The idea is to subsequently conduct two regularization operators, termed *opening* and *closing*, and two compute the difference between the original and both of the regularized version of the binary mask. Formally, the mask manufacturability criterion can hence be defined as:

$$\mu_{MC} := w_o |I_m - o_k(I_m)| + w_c |I_m - c_k(I_m)|, \quad (7.61)$$

where o_k and c_k are the opening and closing operations, respectively. The index k denotes the proximity of the influence of the operation and $|\cdot|$ is a suited norm; the Manhattan norm in this case. The optional weight factors w_o and w_c can be used to attenuated or accentuate any of the two operations. Both operators can be regarded as a composition of two more fundamental morphological operations, *erosion* and *dilation*. In our context, the first operator, erosion can be described as follows: given a pattern consisting of one type of pixels (e.g., bright) and a structuring, for example, a quadratic element. Further let the structuring element be centered on one target pixel in the original pattern. The pixel remains if and only if all pixels in the pattern within the support of the structuring element are identical with the ones in the structuring element. This is illustrated in Figure 7.16 as the first operation in the top and the second operation in the bottom row. To formalize, let A be the set of bright pixels, that is,

$$A := \{\mathbf{z} \in E \mid I_m(\mathbf{z}) = 1\}.$$

Let further B be a structuring element, also defined on E . In our case, it is appropriate to let B define a quadratic region of the grid, i.e.,

$$B_k = \{\mathbf{z} \in \{-k, \dots, k\}^2 \subseteq E\},$$

where $k \in \mathbb{N}$ (the index k is omitted in the following). Erosion can then be defined as follows:

$$A \ominus B := \{\mathbf{z} \in E \mid (\forall \mathbf{b} \in B)(\mathbf{z} + \mathbf{b} \in A)\}. \quad (7.62)$$

Analogously, dilation is an operation that for each pixel in the original pattern reproduces the support of the structuring element, shown in Figure 7.16, where it is the second and first operation in the top and the bottom row, respectively. Formally,

$$A \oplus B := \bigcup_{\mathbf{b} \in B} (A)_{\mathbf{b}}, \quad (7.63)$$

where $(A)_{\mathbf{b}}$ denotes the by \mathbf{b} translated set A .

The opening operator, which can be perceived as a subtle-feature-removing operator, is defined as:

$$A \circ B := (A \ominus B) \oplus B$$

The closing operator, which can be perceived as a subtle-space-removing operator, is defined as:

$$A \bullet B := (A \oplus B) \ominus B$$

The effects of both operations are illustrated in the top and bottom rows of Figure 7.16, respectively. There, a proximity of one pixel ($k : 1$) was used, rendering B a 3×3 matrix representative.




Figure 7.16: Mask manufacturability criterion based on morphological regularity analysis for a 3×3 support ($k : 1$). Top row: Opening consists of an erosion and a dilation operation, removing subtle bright parts from features. Bottom row: Closing, performing dilation and then opening, can be considered the dual operation in that it removes fragile spaces. Image taken from Seifert (2006b).

Aerial Image

After repairing and evaluating the geometry of the mask, the image simulator is called with that mask and the source as inputs. Since in this example no further constraints are applied to the source (as opposed to, for example, the peak intensity minimization explained in Section 7.1.3) the source geometry is simply transcribed into the mesh grid of the simulator and assigned accordingly. Since we are concerned here with an aerial-image-based assessment, no further processing steps are conducted. The aerial image data are evaluated as follows.

The evaluation of lithographic processes has to take several effects into account. Following selection of criteria proved sound in order to achieve well-performing imaging conditions:

CD Criterion The critical dimension (CD) specifies the size of the printed main feature. The deviation between the size of the printed feature at the threshold intensity CD_p and the target CD_t provides the first aerial image (minimization) component of the objective function (cf. Figure 7.17a):

$$\mu_{\Delta CD}(\mathbf{x}) := |CD_p(\mathbf{x}) - CD_t| \quad (7.64)$$

For simple one-dimensional line/space studies, a single CD measurement position is indeed sufficient. In the case of 2-D features, such as contact holes, the CD difference has to be evaluated at different positions. Also Multiple CD evaluations can be combined with a weighted sum approach or using a root mean square error (RMSE). For each evaluation region, an according user-provided cut line has to be defined.




Figure 7.17: Evaluation of lithographic images: (a) main feature: critical dimension (CD), i.e., size of feature at a specified threshold intensity (dashed line), the (average) slope is evaluated at the nominal edges of the feature; (b) global evaluation: band criterion, which ensures that sidelobes are only printed where required.

Band Criterion (BC) In order to ensure that only the desired features are printed, those solutions whose sidelobes exhibit a low intensity are punished. On the other hand, the feature's intensity should not exceed a certain value to avoid exposed spots within areas that are to be printed. For that, a band around the threshold value (as shown in Figure 7.17b) has been introduced. Solutions that violate this criterion are punished.

$$\mu_{BC}(\mathbf{x}) := \begin{cases} I_t(1 + p_t) - I_e(\mathbf{x}) & \text{no-print sidelobes with too low an intensity} \\ I_e(\mathbf{x}) - I_t(1 + p_t) & \text{print areas with too high an intensity} \end{cases} \quad (7.65)$$

where I_t is the threshold intensity, I_e is the intensity at the minimum of the sidelobe (maximum, resp.), and p_t is the band size (which is set to $\pm 30\%$ throughout this work). Again, multiple features or orientations can be combined by introducing several cut lines and by combining the assessment result using a weighted sum or RMSE approach.

Slope Criterion (SC) This criterion (see Section 2.2.7 for a discussion) ensures that good solutions should be tolerant against variations in exposure doses. Thus, the slope of the normalized intensity curve at the edges of the target feature is to be maximized (cf. Figure 7.17a).

$$\mu_{SC}(\mathbf{x}) := dI/dx|_{\text{edge}} \quad (7.66)$$

Through-focus evaluation In order to provide a stable process, variations of the focus position should have as little impact on the imaging performance as possible. To optimize the through-focus behavior, the former criteria are conducted at a number of different focus positions.

Fitness Function

The combination of the individual criteria is critical in that it may drastically impact both the result and the convergence rate of the optimization process. Especially, the determination of a suited balance between the individual penalty and merit values is far from being straightforward since it can easily lead to an unintentional bias (see Section 4.4 and Section 4.11 for a discussion). In a first step, using

the single-objective genetic algorithm, we have used a weighted sum approach. For the combined fitness function, we thus have that,

$$\mu : \mathbf{x} \mapsto \frac{w_{MC} \cdot \mu_{MC}(\mathbf{x}) + w_{AI} \cdot \sum_{i=1}^n \frac{w_{\Delta CD} \cdot \mu_{\Delta CD}(\mathbf{x}) + w_{SC} \cdot \mu_{SC}(\mathbf{x}) + w_{BC} \cdot \mu_{BC}(\mathbf{x})}{n(w_{\Delta CD} + w_{SC} + w_{BC})}}{w_{MC} + w_{AI}} \quad (7.67)$$

where μ_{MC} , $\mu_{\Delta CD}$, μ_{SC} , and μ_{BC} denote the results of the former criteria and w_{MC} , $w_{\Delta CD}$, w_{SC} , and w_{BC} symbolize the corresponding weights, w_{AI} is used to bias the ratio between geometry (mask) and aerial image evaluation. The number of focus settings, which each solution is evaluated for, is denoted by n .

To exploit the multi-objective nature of the problem and to circumvent the problems associated with a weighted-sum replacement problem, we have applied a multi-objective GA to a number of examples. As a consequence, the fitness function also becomes multi-objective, that is, results a vector rather than a scalar. In fact, one could employ as many sub-objectives as there are criteria, which, however, is not practical for a number of reasons; for example:

- (1) The solution space after optimization becomes intractable, consisting of too many candidates for the decision maker to select from.
- (2) Due to the “curse of dimensionality,”¹⁸ genetic operations are significantly impaired and thus lead to a poor performance.
- (3) Especially crowding and clustering techniques in the objective space tend to fail since they are less likely to reject or combine solutions.¹⁹

To reduce the number of sub-objectives, we have combined the criteria as follows (compare Seifert, 2006b). For the CD different criterion:

$$\mu_{\Delta CD} : \mathbf{x} \mapsto w_0 \mu_{\Delta CD}^0 + \sum_{\phi} \mu_{\Delta CD}^{\phi}, \quad (7.68)$$

where ϕ is the focus position and $\mu_{\Delta CD}^{phi}$ is the aggregate CD difference at the cut lines at that focus position. In order to accentuate the best-focus performance, a weight w_0 can be optionally used. Similarly for the band criterion,

$$\mu_{BC} : \mathbf{x} \mapsto w_0 \mu_{BC}^0 + \sum_{\phi} \mu_{BC}^{\phi}, \quad (7.69)$$

and the slope criterion,

$$\mu_{SC} : \mathbf{x} \mapsto w_0 \mu_{SC}^0 + \sum_{\phi} \mu_{SC}^{\phi}, \quad (7.70)$$

where the same weights w_0 as in the CD case can be used.

7.3 Results

The optimization courses presented in this section have been conducted using a light’s wavelength (λ) of 193 nm and a numerical aperture (NA) of 0.7.²⁰

7.3.1 One-dimensional Mask Representation

In a first experiment, the objective of the optimization is to find the optimum mask for parameters for a lines/spaces pattern with a line width of 110 nm and pitch between the lines of 760 nm. A reference process window (exposure dose vs. defocus) is defined by a dense lines and spaces pattern (same line width, a pitch of 220 nm) with a threshold intensity of 0.314. This experiment was performed using the

¹⁸A term first coined by Bellman (1961).

¹⁹For examples, see Silverman (1986).

²⁰Parts of this section have been previously published in (Fühner et al., 2007a).

continuous mask representation: Three assist features have been placed between the two main features, one of which is placed in the center between both features, the remaining two assists have identical sizes and the same distance from the main feature. The goal was to find optimal width and distance parameters as depicted in Figure 7.18a. The ranges for the four parameters are set as follows:

- width of main feature: $w_m \in [40 \text{ nm}, 100 \text{ nm}]$,
- width of center assist $w_c \in [40 \text{ nm}, 80 \text{ nm}]$,
- width of the other two assists $w_a \in [40 \text{ nm}, 80 \text{ nm}]$, and
- distance of these assists $d_a \in [40 \text{ nm}, 300 \text{ nm}]$.

width of main feature: $w_m \in [40 \text{ nm}, 100 \text{ nm}]$, width of center assist $w_c \in [40 \text{ nm}, 80 \text{ nm}]$, width of the other two assists $w_a \in [40 \text{ nm}, 80 \text{ nm}]$, and distance of these assists $d_a \in [40 \text{ nm}, 300 \text{ nm}]$. For this example, the illumination parameters are not varied but fixed to $\sigma_{\text{in}} : 0.5$, $\sigma_{\text{out}} : 0.7$.

Table 7.1: GA settings for both the continuous and discrete mask representation. Five runs are performed for each setting. The window size specifies the number of compared individuals with restricted tournament selection (see Section 4.9). The maximum, minimum, and average fitness of the five runs' best solution are listed.

selection operator	window size	mutation rate	Continuous			Discrete		
			average fitness	lowest fitness	best fitness	average fitness	lowest fitness	best fitness
roulette wheel		0.001	55.58	55.33	56.42	28.24	26.94	28.88
roulette wheel		0.005	55.93	55.35	56.42	28.60	26.51	29.80
binary tournament		0.001	55.52	55.35	55.67	24.37	22.84	25.49
binary tournament		0.005	55.60	54.78	56.42	29.45	28.52	29.80
restricted tournament	20	0.001	56.16	55.60	56.42	28.91	28.10	29.80
restricted tournament	20	0.005	<u>56.42</u>	<u>56.42</u>	<u>56.42</u>	29.68	28.99	<u>30.08</u>
restricted tournament	50	0.001	<u>56.42</u>	<u>56.42</u>	<u>56.42</u>	29.64	<u>29.01</u>	29.80
restricted tournament	50	0.005	<u>56.35</u>	<u>56.11</u>	56.42	29.57	28.77	29.80

The number of individuals is 51; experiments with larger population numbers did not improve the convergence behavior—yet, in case of restricted tournament selection, a remarkable slowdown could be observed. Crossover probability is set to 0.6, mutation probabilities are 0.001 and 0.005, and the number of generations is 500. As shown in Table 7.1, several settings are tested. In order to verify the reproducibility of results, five runs are conducted for each setting; each run takes about 5 – 6 minutes on a five node cluster with each machine having two 2.66 GHz Pentium 4 processors. Although with all settings the GA results do not differ significantly, it is noticeable that RTS provided the highest degree of reproducibility of results.

The improvement of the imaging performance after a certain number of generation is depicted in Figure 7.18b. In the first row, aerial images at a defocus of 250 nm are presented. The dashed line indicates the threshold intensity which is used for the image evaluation. The lower row shows process windows for the optimized mask (solid lines) compared to dense lines/spaces (dashed lines). On the left, simulation results for the mask with no assist features are shown. The minimum intensity almost reaches the threshold value. Also, the overlap between the process windows of semi-dense and dense patterns is very small, i.e., semi-dense and dense patterns can only be printed simultaneously within a small range of dose and focus settings. After 50 generations the GA has already provided mask patterns that result in a highly improved imaging performance. Both quality and overlap are noticeably improved.

The same problem is approached with the discrete mask representation. As illustrated in Figure 7.19, in the first generations the GA is assembling clusters of lines. In the 10th generation the mask mainly consists of thin lines and spaces. However, after a number of iterations the manufacturability criterion leads to larger regions. In the following steps, the GA combines these clustered patterns. As the clus-

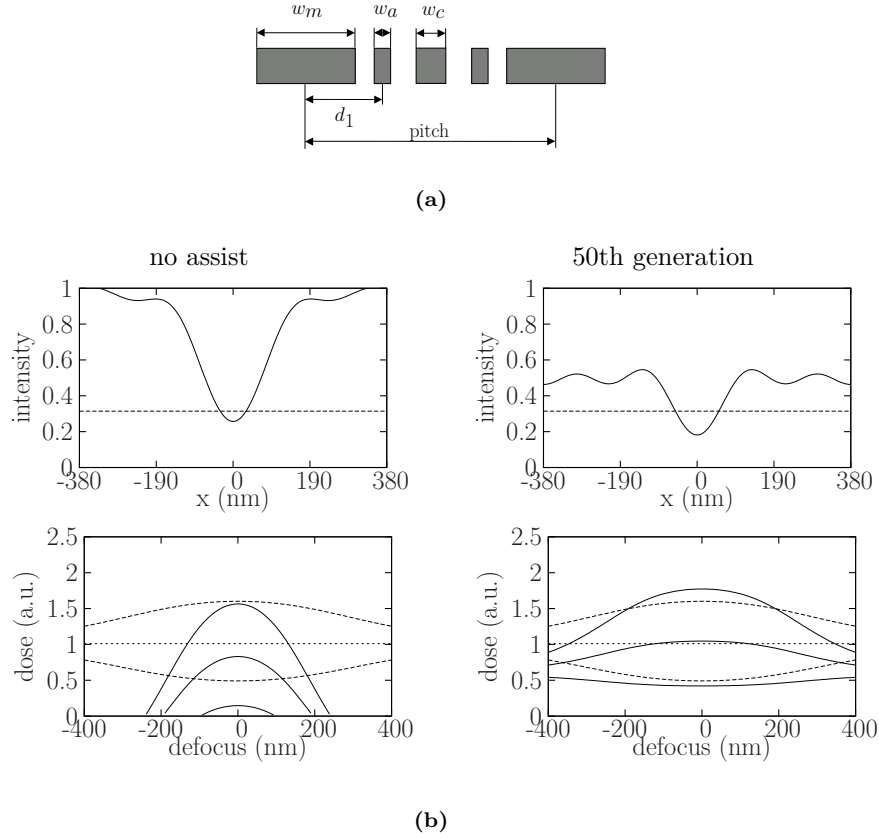


Figure 7.18: Results for one-dimensional mask representation. (a) Mask geometry with 4 variables; (b) Improvement of the image performance obtained with the continuous mask representation, upper row: images at 250 nm defocus (solid lines) and threshold intensity (dotted line), lower row: process windows of the semi-dense feature (solid lines, target: width = 110 nm, pitch = 760 nm) compared to process windows of dense features (dashed lines, target: width = 110 nm, pitch=760 nm), left column: isolated line without assist ($w_m = 150$ nm), right column: after 50 generations ($[w_m, w_a, w_c, d_1] = [125$ nm, 50 nm, 55 nm, 215 nm]).

tering phase requires some time, the number of generations has to be increased. In this example, the GA was run with 2000 generations, computing time (using the same computer set-up as in the previous example) scaled up to about 45 minutes. After 800 generations, the GA yielded a mask geometry which guarantees a good imaging performance and which resembles the mask obtained from the continuous optimization approach. The different settings tested with this experiment are listed in Table 7.1.

The discrete representation of mask geometries can be used to compare different mask types and illumination options. This is demonstrated in Figure 7.20, which shows the best mask layouts, annular illumination settings, and resulting process windows. The only information the GA was provided with was the size of the addressable grid on the mask ($\Delta x = 5$ nm), the complex transmission values of the mask (binary: $1 + 0j$, $0 + 0j$; PSM: $1 + 0j$, $-1 + 0j$, $0 + 0j$), and the type of illumination (annular). In case of binary masks, the optimization of the merit function 7.67 results in a typical sub-resolution assisted mask configuration and strong off-axis illumination. For the PSM, the algorithm suggests a standard phase edge design and coherent illumination. The process window for the PSM is considerably larger than for binary masks.

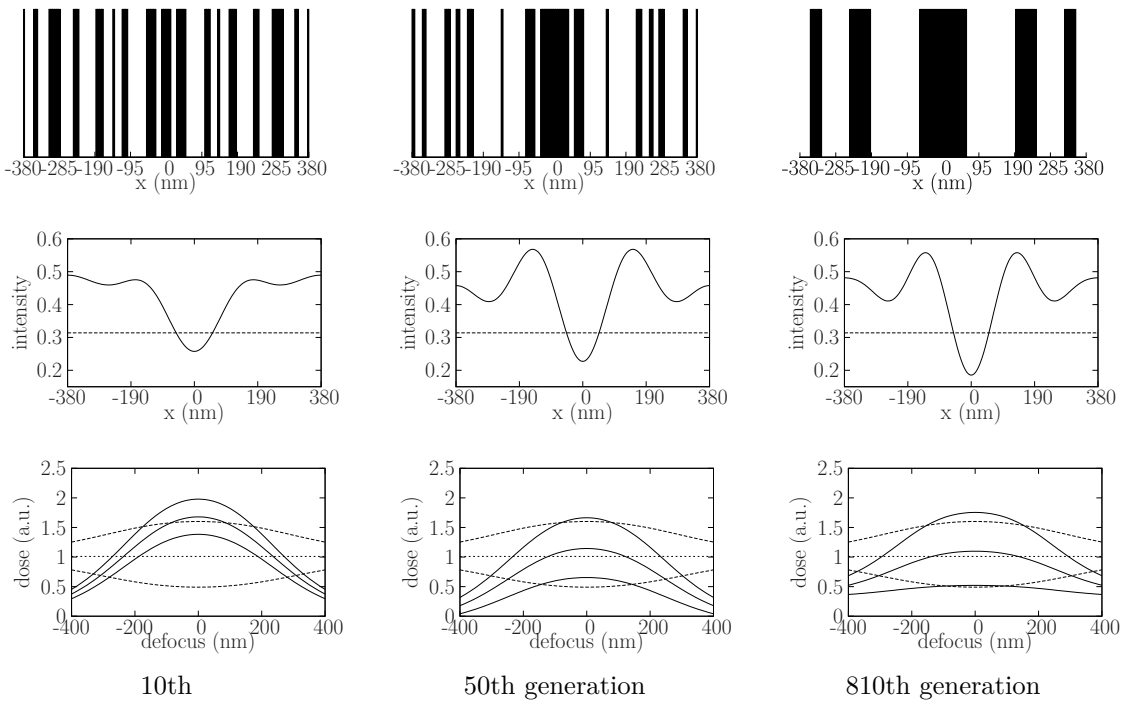


Figure 7.19: Improvement of the image performance obtained with the GA for the discrete mask representation. Upper row: mask geometry, center row: images at 250 nm defocus, lower row: process windows.

Proximity Curves for Lines/Spaces Patterns

In general, among other structures, a lithographic mask typically contains lines and spaces patterns with different pitches (periods). These structures have to be printed with an identical illumination set-up. Thus the optimum mask configuration can be determined as function of the pitch.

Figure 7.21 presents the result of such an optimization procedure, using the 1-D continuous mask representation. For that, the pitch was varied in a range of 220 nm (dense features) to 1000 nm (semi-dense features) using 20 nm steps. For small pitches (up to 320 nm), the space between main features does not allow for assists to be placed. At pitches starting from 340 nm, a single assist between two main features provides the best lithographic performance. As the pitch is increased, more additional assists have to be placed between the main features for an ideal imaging performance. At the largest pitch of 1000 nm, four assists between the two main features provide the best solution.

The presented results also demonstrate the “forbidden” pitch phenomenon (Chen et al., 1997; Mansfield et al., 2000; Trouiller et al., 2002): At pitch widths close to the introduction of an additional assist, a poor imaging performance (fitness) can be observed. This information can be retrieved from the optimization results without any additional a priori information of the corresponding process.

7.3.2 Two-Dimensional Mask Representation

The following example demonstrates a first application of the two dimensional mask representation. The goal was to find optimum mask and illumination settings for contact holes (140 nm · 170 nm). The following pitch settings have been investigated:

- (a) dense: x -pitch = 320 nm, y -pitch = 410 nm,
- (b) chain: x -pitch = 320 nm, and
- (c) isolated.




Figure 7.20: Optimized imaging performance for 90-nm isolated lines for the discrete mask representation, upper row: binary mask (BIM), lower row: alternating phase shift mask.

For the source geometry, the parametric representation (cf. Section 7.2.1) has been used. As the principal set-up for this representation is fixed, several different layouts have been used. Illumination geometries with their corresponding parameters are listed in Table 7.2.

Table 7.2: Illumination source layouts and corresponding parameters (cf. Figure 7.11).

configuration	parameter values
CIRCULAR	$0.3 \leq \sigma \leq 0.9$
ANNULAR	$\sigma_{\text{in}} \geq 0.3$ $\sigma_{\text{out}} - \sigma_{\text{in}} \geq 0.12$
MP0 (0-deg. quadrupole)	$\theta : 0 \text{ deg}$ $0.45 \leq \sigma_{\text{off}} \leq 0.75$ $0.12 \leq \sigma_r \leq 0.4$
MP45 (45-deg. quadrupole)	$\theta : 45 \text{ deg}$ $0.45 \leq \sigma_{\text{off}} \leq 0.75$ $0.12 \leq \sigma_r \leq 0.4$
DP0 (0-deg. dipole)	$\theta : 0 \text{ deg}$ $0.45 \leq \sigma_{\text{off}} \leq 0.75$ $0.12 \leq \sigma_r \leq 0.4$
DP90 (90-deg. dipole)	$\theta : 90 \text{ deg}$ $0.45 \leq \sigma_{\text{off}} \leq 0.75$ $0.12 \leq \sigma_r \leq 0.4$

Optimization runs were performed for AttPSMs with both high (19.3 percent) and low (6 percent) background transmission values. As in the one-dimensional case, the optical settings are: numerical aperture (NA) of 0.7 and a wavelength (λ) of 193 nm. The threshold for the aerial image evaluation was computed using a 160-nm dense lines and spaces pattern with a 320-nm pitch. For the band criterion (cf. Section 7.2.2), a threshold band of 30 percent was introduced. That is, main features are assumed to print only if their intensity exceeds the threshold intensity by at least 30 percent. Vice versa, sidelobes

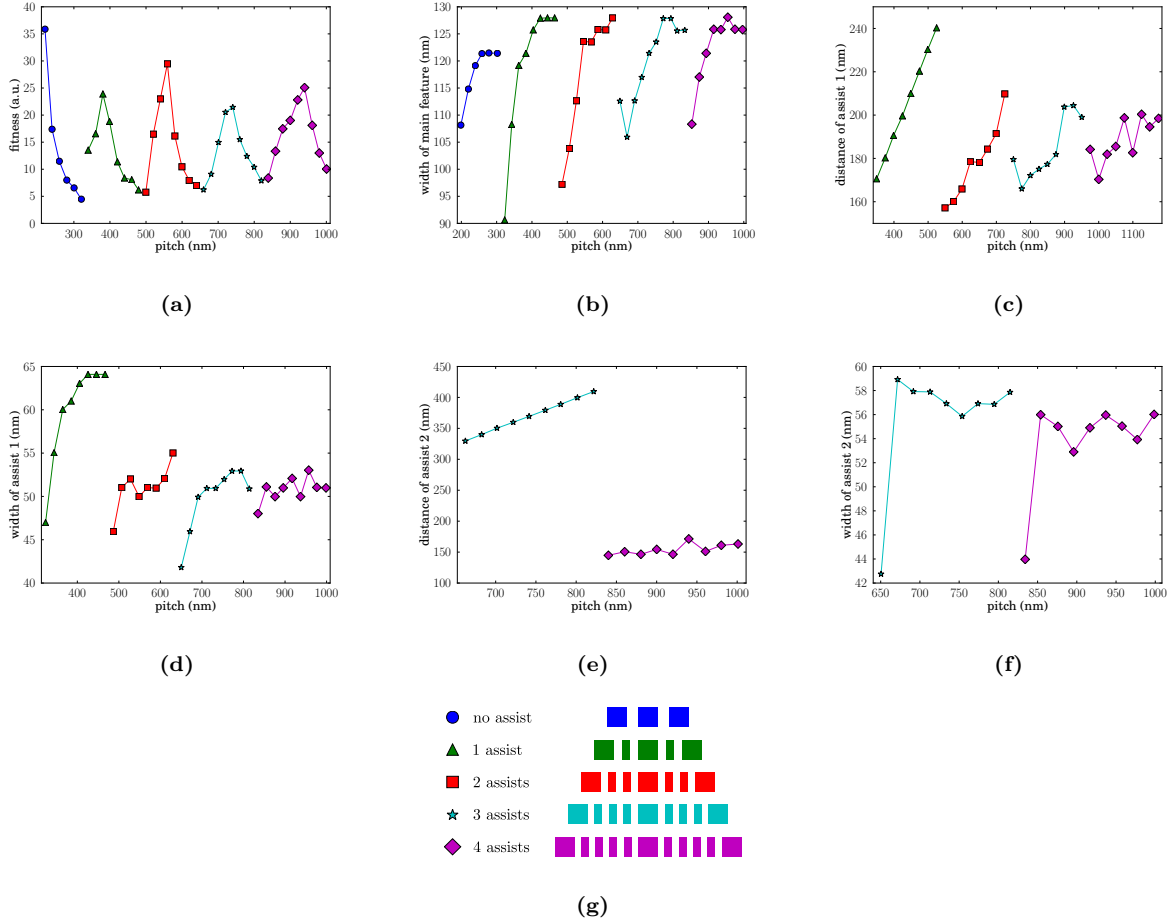

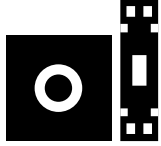

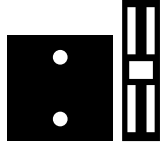
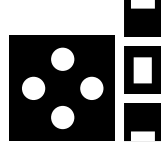
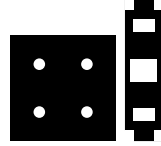
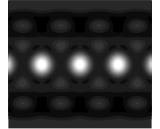
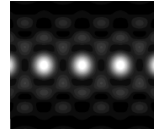
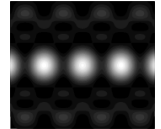
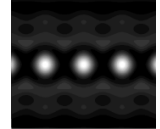
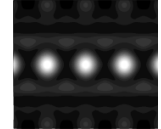
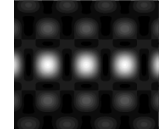
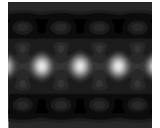
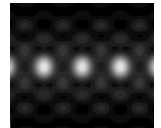
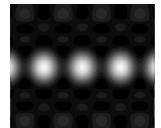
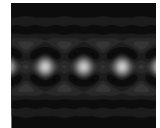
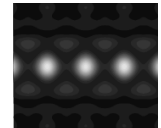
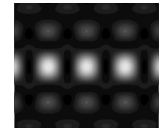
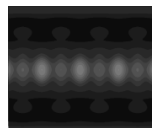
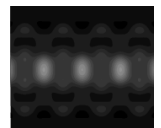
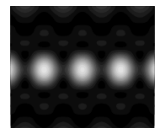

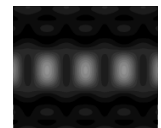
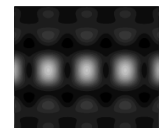


Figure 7.21: Simulated optimum main feature and scattering bar number, size and placement configuration for 110-nm lines with pitches ranging from 220 nm to 1000 nm. The optimal parameters of the annular illumination were found to be $\sigma_{in}/\sigma_{in} = 0.51/0.72$, all other parameters as specified before: (a) Maximum fitness per pitch. (b) Optimum main feature width per pitch, (c) Optimum distance of first assist from main feature (center-to-center). In the case of only one assist (green triangles), the assist is placed in the center of two adjacent main features. (d) Optimum width of the first assist. (e) Optimum distance of the second assist from main feature. In the case of three assists (cyan-colored stars), the second assist is placed in the center of two adjacent main features. (f) Optimum width of the second assist. (g) Assist configuration schemes.

Table 7.3: Results for different illumination settings for a chain set-up using low transmission (6percent) attenuated phase shifting mask. The first row shows the mask layout (right) and the illumination type (left). The second row displays the fitness (μ). And the images and the respective CDs obtained at no, 200 nm, and 400 nm defocus are depicted in the last three rows.






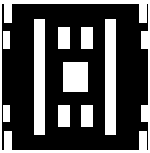
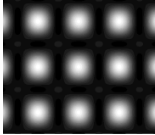
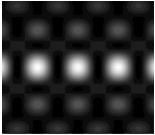
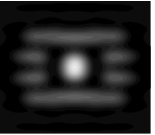
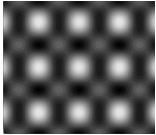
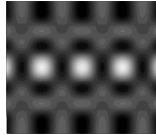
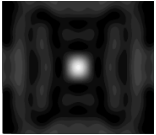
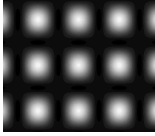
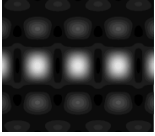

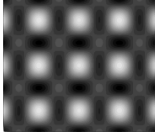
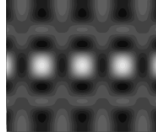

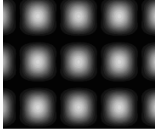
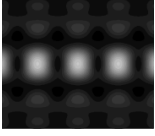
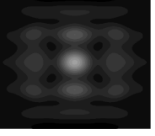

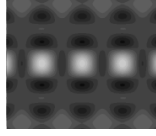

CIRCULAR	ANNULAR	DP0	DP90	MP0	MP45
					
μ : 123.6	μ : 161.7	μ : 514.6	μ : 87.0	μ : 112.6	μ : 401.0
					
CD _x : 140.1 nm CD _y : 170.0 nm	CD _x : 140.0 nm CD _y : 170.0 nm	CD _x : 140.1 nm CD _y : 170.1 nm	CD _x : 140.1 nm CD _y : 170.0 nm	CD _x : 140.3 nm CD _y : 170.2 nm	CD _x : 140.0 nm CD _y : 170.0 nm
					
CD _x : 123.0 nm CD _y : 154.8 nm	CD _x : 124.7 nm CD _y : 161.4 nm	CD _x : 136.1 nm CD _y : 170.2 nm	CD _x : 124.0 nm CD _y : 141.1 nm	CD _x : 124.1 nm CD _y : 155.6 nm	CD _x : 132.2 nm CD _y : 169.1 nm
					
no CD _x no CD _y	CD _x : 45.5 nm CD _y : 88.1 nm	CD _x : 126.5 nm CD _y : 165.7 nm	no CD _x no CD _y	no CD _x no CD _y	CD _x : 108.7 nm CD _y : 159.9 nm

will print up to an intensity of 30 percent below the threshold. Structures were assumed to print at 30 percent above this threshold (and not to print at 30 percent below).

Table 7.3 shows results for the low transmission AttPSM and a *chain* (periodic in x direction) set-up. Although for all illumination layouts good results are achieved for 0 nm and even 200 nm defocus positions, only the dipole and quadrupole geometries show a stable imaging behavior for the 400-nm defocus case. In all cases, the proposed optimization procedure is able to place assist structures and alter the illumination geometry parameters such that a drastically improved imaging performance is obtained.

Extensive optimization runs have been conducted using the 45-degree quadrupole illumination geometry (MP45). Both low and high transmission AttPSM have been investigated. The results are depicted in Table 7.4. For both mask types and all set-ups well-performing mask and illumination settings could be obtained. Even though the background for the high transmission masks is naturally brighter, no sidelobe problems were encountered.

Table 7.4: Result of the mask/source optimization for dense, chain, and isolated 140 nm · 170 nm contact holes with MP45 illumination for: low transmission (6 percent) and high transmission (19.3 percent) attenuated PSM. The upper row show a unit cell of the optimized mask layout. The lower three rows show images and CDs at nominal focus, 200 nm and 400 nm defocus.

low transmission (6 percent)			high transmission (19.3 percent)		
dense	chain	isolated	dense	chain	isolated
					
					
CD _x : 140.0 nm CD _y : 170.0 nm	CD _x : 140.0 nm CD _y : 170.0 nm	CD _x : 140.0 nm CD _y : 170.0 nm	CD _x : 140.0 nm CD _y : 170.0 nm	CD _x : 140.0 nm CD _y : 170.0 nm	CD _x : 140.0 nm CD _y : 170.0 nm
					
CD _x : 134.3 nm CD _y : 167.3 nm	CD _x : 132.2 nm CD _y : 169.1 nm	CD _x : 130.4 nm CD _y : 148.4 nm	CD _x : 140.0 nm CD _y : 171.1 nm	CD _x : 138.9 nm CD _y : 169.2 nm	CD _x : 132.4 nm CD _y : 160.2 nm
					
CD _x : 116.4 nm CD _y : 156.3 nm	CD _x : 108.7 nm CD _y : 159.9 nm	CD _x : 56.7 nm CD _y : 53.7 nm	CD _x : 136.9 nm CD _y : 164.6 nm	CD _x : 131.8 nm CD _y : 160.6 nm	CD _x : 101.3 nm CD _y : 121.7 nm

7.3.3 Track/Sector Illumination Representation

The following results were obtained with the more flexible track and sector-based illumination representation (see Section 7.2.1). Additionally, computation time for the simulation and evaluation of the aerial image was reduced by introducing a pre-evaluation step: Before the actual evaluation is conducted, a coarse grained (fewer mesh points) simulation—at the best focus position only—is performed. Only if the resulted CD of the aerial image is close enough to the target CD, a computation on a finer grid and at all defocus positions is performed. Thus it was possible to remarkably improve the convergence behavior and hence the obtained solutions.

Figure 7.22 shows the best result for a dense set-up. The illumination was divided into seven sections (defined in a quarter-circle) and ten tracks. In order to improve the convergence time, four inner tracks have been set to 0 transmission (forbidden tracks). Acceptable results at 200 nm defocus were obtained after about 1500 generations (300 solutions per generation), which corresponds to about 4 hours (on 20 CPUs; 2.8 GHz Pentium 4). In this case, the following CD mismatch values could be obtained:

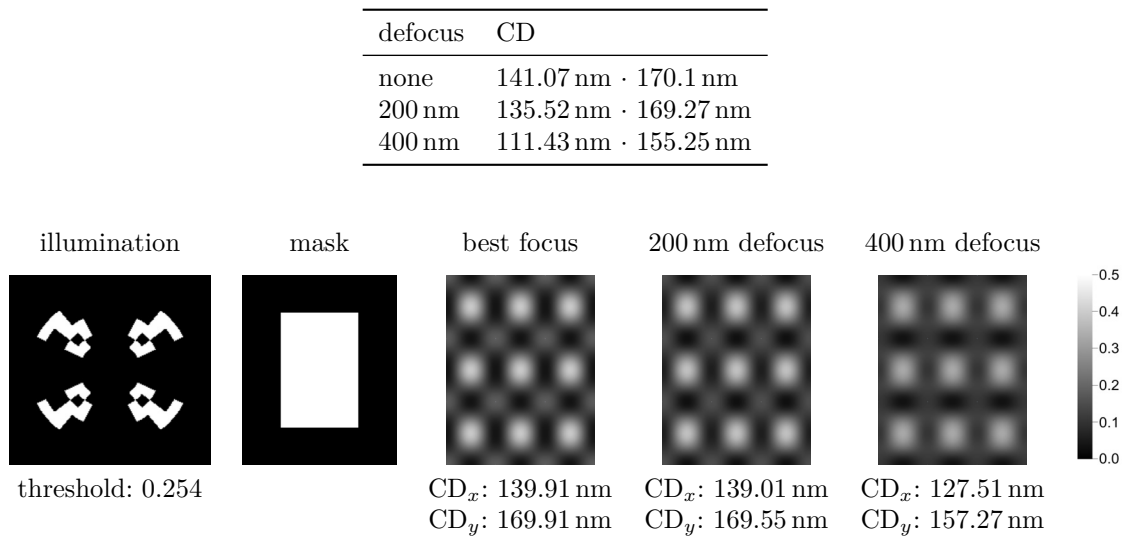


Figure 7.22: Best optimization results achieved for a dense contact hole set-up.

The result shown in Figure 7.22 was achieved after a prolonged run of 7000 generations (\equiv 16 hours). Although the optimizer was allowed to place as many as five rectangles in each quadrant, no assist features are generated. The obtained solution provides an excellent CD match for the best focus but also for the 200 nm defocus position. At 400 nm the CD is still close to the target. None of the defocus positions exhibit sidelobe printing.

The best optimization results of all feature settings were obtained for a chained set-up (see Figure 7.23). Even at 400 nm defocus the CD is on target. Again, at no defocus position does the solution run the risk of sidelobe printing. As in the former example, the illumination consists of seven sections and ten tracks, with four forbidden tracks. After 300 generations (600 solutions/generation; \equiv 1.5 h on 50 CPUs) an acceptable solution was acquired:

defocus	CD
none	143.11 nm · 169.5 nm
200 nm	141.39 nm · 169.13 nm
400 nm	135.68 nm · 169.1 nm

After about 1000 generations (4 h), the following results were obtained:

defocus	CD
none	140.38 nm · 170.03 nm
200 nm	139.82 nm · 169.78 nm
400 nm	140.32 nm · 169.94 nm

There, even in the 400-nm case, the CD values are almost exactly on target.

The illustrated solution (Figure 7.23) is resulted after a long-lasting run of 7000 generations (\equiv 24 h). Compared to the results after 1000 iterations, only slight improvements of the best focus' CD and the exposure latitude were achieved.

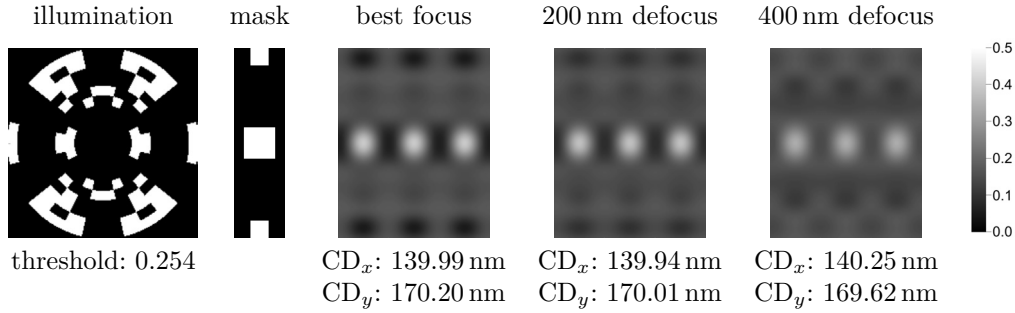


Figure 7.23: Best optimization results achieved for a chained contact hole set-up.

Figure 7.24 illustrates the best optimization result for an isolated contact hole. The solution was achieved after 3500 generations (500 individuals/generation; 28 h on 20 CPUs). The illumination was pixelated as in the former examples. The imaging performance is slightly worse than for the previous set-ups. However, the target CD was reached at the best focus position, and also at 200 nm defocus it is almost on target. For 400 nm defocus, the disagreement is significant. However, at any position no sidelobe printing appears.

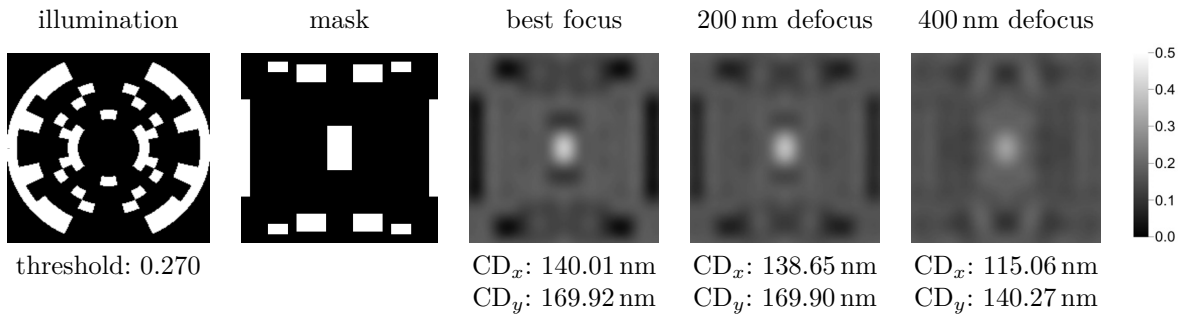


Figure 7.24: Best optimization results achieved for an isolated contact hole.

After 500 generations (\equiv 5 h), the following decent CD values could be obtained:

defocus	CD
none	141.12 nm · 169.67 nm
200 nm	133.84 nm · 170.3 nm
400 nm	90.39 nm · 139.33 nm

The following example (Figure 7.25) is a first demonstration of a mutual optimization of two different feature set-ups. In this case, an isolated contact hole and a chained set-up have been simultaneously optimized. As an additional objective, a source illumination, ideal for both masks, had to be found. A first

acceptable image performance could be obtained after 1000 generations (400 individuals/generation; 20 h on 20 CPUs):

defocus	CD isolated	CD chain
none	148.62 nm · 172.12 nm	162.33 nm · 172.46 nm
200 nm	145.59 nm · 172.28 nm	159.25 nm · 172.41 nm
400 nm	124.67 nm · 158.6 nm	143.06 nm · 162.84 nm

Results depicted in Figure 7.25 were acquired after 4000 generations, or about 72 hours.

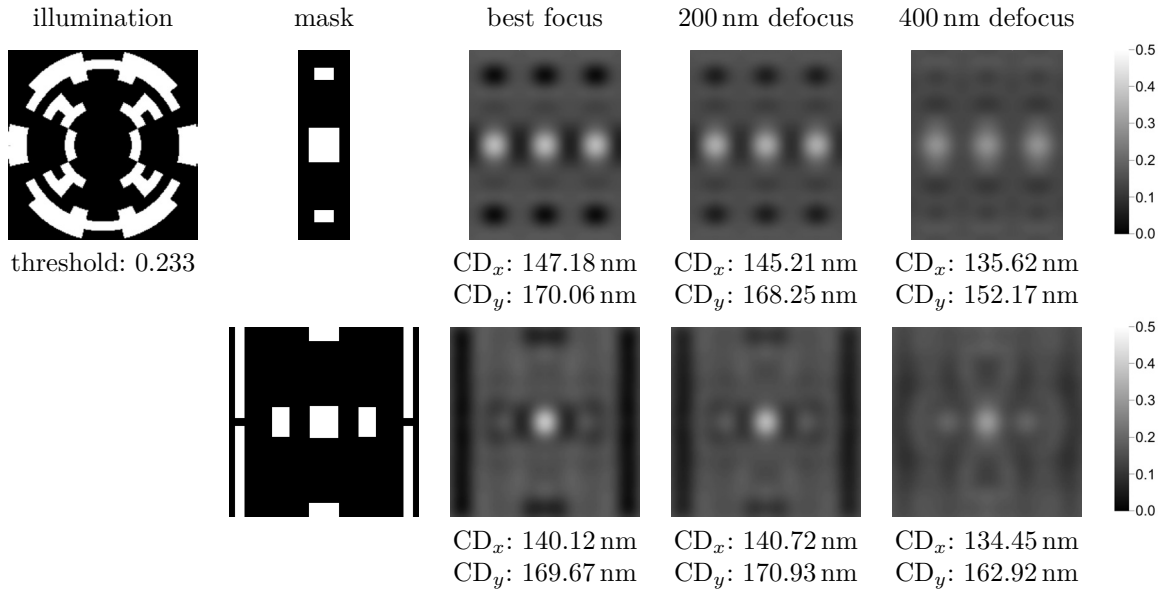


Figure 7.25: Best result of a mutual optimization for both chain set-up (top) and an isolated contact hole (bottom).

The computation time for this first multi-feature example is still very high. This is due to a sloppy formulation: the CD criterion of the merit function consists of the average CD mismatch of the two features. Thus a bad CD match of one feature and an excellent merit of the other set-up results in a similar fitness as sound CD performances of both features. In order to stabilize the convergence behavior, a poor fitness behavior of a single feature has to be “punished” more drastically.

7.3.4 Results Using the Multi-objective GA

The following results were obtained using SPEA2 using an unrestricted archive. Both the parameter representation and the aerial image evaluation were adopted from the single-optimization approach. For the manufacturability criterion, first, the original procedure was inherited, but was later replaced by the dilation/erosion operator. In contrast to the single-objective optimization approach, the formulation of the actual fitness function was greatly simplified: no weights for individual sub-criteria (Δ CD, band, slope, etc. criteria) had to be determined.

As with this multi-objective approach, the GA converges towards the Pareto front (set of equally well performing solutions), it also yields candidate solutions that are performing well in one trait but abysmally in other objectives. Therefore a post-optimization review and selection among the set of solutions with their trade-offs is required. Figure 7.26 is a rather extreme example of a non-viable candidate solutions for the isolated contact hole problem. This example exhibits an inverse intensity situation: The intensity of the main feature is low, its surroundings show a high intensity. Still, the

CD is on target, and hence the solution receives a very high fitness for the CD criterion. However, the band criterion is obviously strongly violated. Because of its high fitness regarding one dimension (CD criterion), this solution is situated on the Pareto front approximation.




Figure 7.26: Example of a very poorly performing solution. The band criterion's violation is extreme, but the solution has very good CD values (albeit measured at different positions outside the target CD area).

Figure 7.27 illustrates such Pareto approximations, taken from a multi-dimensional objective space and projected onto the plane of the band and the CD criterion. Review of the lower left area yields reasonable solutions. The structure of the approximation set is an indicator of possible trade-offs, allowing design space exploration to take place.




Figure 7.27: An approximation set found during a run with 200 individuals at generation 1000. The band criterion is drawn in horizontal, the CD criterion in vertical direction. The individual of Figure 7.26 is located in the far lower right part.

Avoiding the intricate definition of weights for the fitness function and using CD deviation, band criterion violation, and manufacturability as direct objectives (and manual post-selection of solutions), we were easily able to obtain well performing results for a chain of contact holes (see Figure 7.28).

As candidate solution were observed to be predominantly situated in specific regions of the Pareto front approximation, additional diversity maintaining measures were incorporated into the multi-objective GA (Section 4.11). Additionally, the improved dilation/erosion manufacturability criterion was introduced. With that, excellent results could be produced for the isolated contact hole set-up. After about 42 h of computation on 8 AMD Opteron 64 cluster nodes, using 250 individuals over 1000 generations, well performing solutions were obtained. A candidate manually selected by review (of the promising lower left area of the Pareto front approximation) is shown in Figure 7.29.

In a next step, the multi-objective algorithm has been applied to the mutual optimization of different feature set-ups. The intuitive choice of an objective vector—the concatenation of the isolated and the

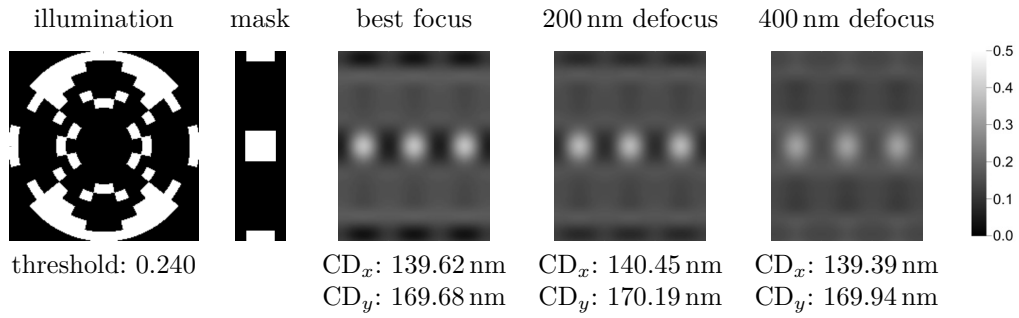


Figure 7.28: A good solution manually chosen from a Pareto front approximation of the contact hole chain set-up. Found after 2000 generations of 400 individuals. Comparably good results were, however, already obtained after 1000 generations of 200 individuals.

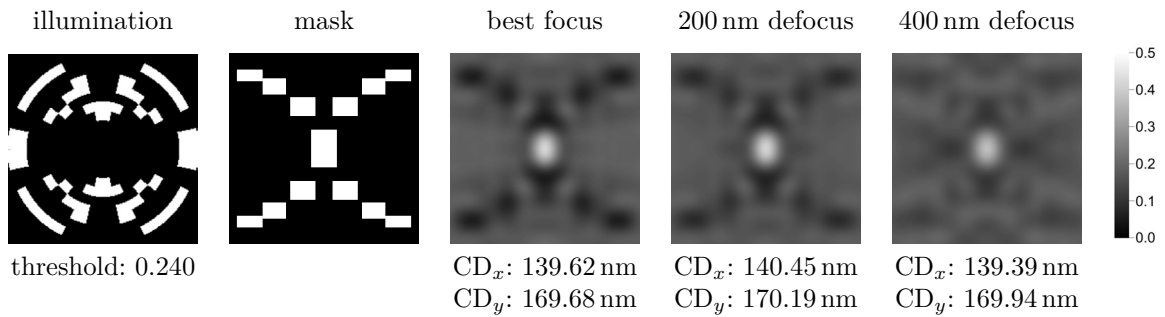


Figure 7.29: A good solution for the isolated set-up.

chain set-ups' criteria evaluation—lead to a very slow convergence and yielded only unacceptable, local optimum solutions. Reducing the dimensionality of the objective space by adding the CD and band criteria of both set-ups yielded a significantly improved convergence. A run of 300 individuals over 2000 generations produced promising results. A typical solution is shown in Figure 7.30. Early convergence, however, is still an issue and better results are expected to be possible (e.g., by introducing further diversity maintaining measures).

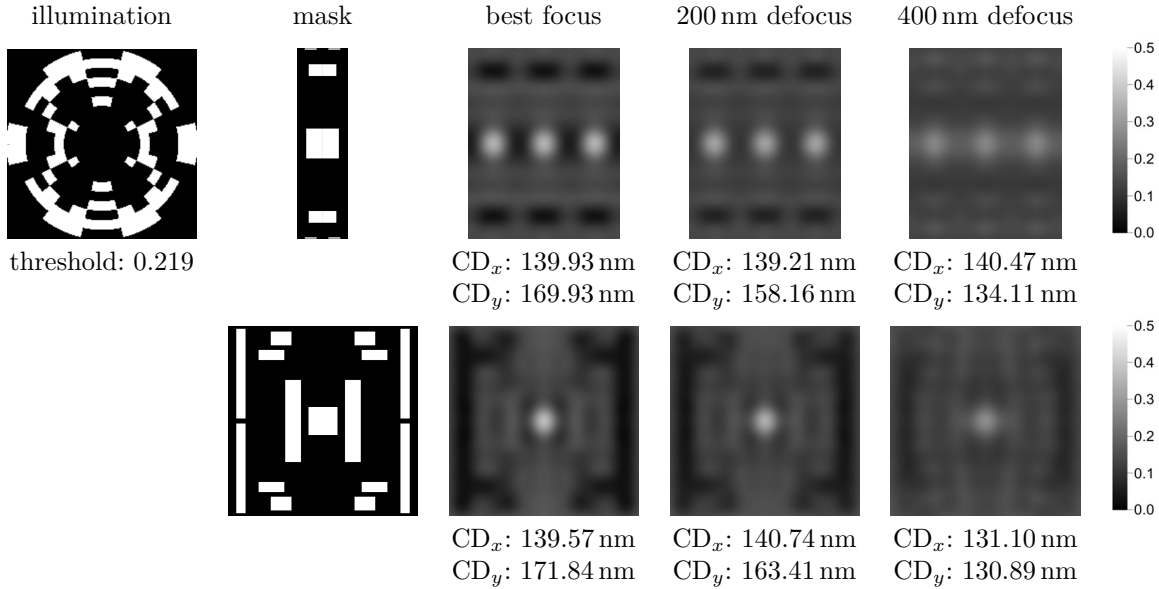


Figure 7.30: A good solution of simultaneous optimization of chained and isolated set-up.

7.4 Source/Mask/Projector Pupil Optimization

The previous examples²¹ clearly demonstrate that a rigorous source/mask optimization approach is generally capable of extending the applicability of lithography generations far beyond their expected lifetime. Through the following numerical experiments we shall underline this argument by allowing for an additional degree of freedom, the projector.²² Aiming at state-of-the-art hyper-NA set-ups, however, additional difficulties have to be accounted for. For example, among other topological mask effects, phase variations may lead to a pronounced asymmetry of process windows (Erdmann, 2001; Ruoff et al., 2008; Erdmann et al., 2010), more severely yet, heterogeneously affecting different features and pitches (Brunner, 1997; Flagello et al., 1997) and thus prohibiting a global compensation. Recently, it has been proposed to modify the projector to incorporate aberrations that are suitable to counteract the mask-induced wavefront skews (Finders et al., 2010; Evanschitzky et al., 2011; Sears et al., 2011). As a number of modern scanner tools already employ a sophisticated wavefront control mechanism, this technique can be expected to be principally feasible, and in fact, has already been demonstrated by Finders et al. (2010). The following experiment is aimed at a co-optimization of the source, the mask and the projector wavefront.

²¹and related work in the field

²²The following results have been presented in (Fühner et al., 2012).

7.4.1 Problem Definition

The example, adopted from Finders et al. (2010)—and slightly modified—consists of a set of different pitches of horizontally and vertically oriented features with a linewidth of 45 and 90 nm, respectively, which are to be patterned in one process. The goal is hence to maximize the common process window. Although the maximization of process windows was inherently targeted also in the former examples, by maximizing the local contrast and the through-focus-behavior, here we have stated it as a more explicitly formulation as will be shown below. Exhibiting a k_1 factor of about 0.3, this example is close to the limit of what is considered the comfort zone of i193 (immersion lithography at a wavelength of 193 nm and an NA of 1.35) printing. As the mask technology, we employ a COG reticle. In contrast to the former examples, the topology of the mask is considered, and the absorber is assumed to be 80 nm thick. The following pitches are considered: 90, 125, 135, 270, 315 nm for the vertical and 207 and 270 nm for the horizontal lines (see Figure 7.31).




Figure 7.31: Seven line/space configurations are to be simultaneously optimized: five different pitches of horizontally oriented 45-nm lines and two pitches of a horizontal 90-nm line.

The mask diffraction spectra are computed through two-dimensional rigorous EMF simulations, using the waveguide method (see Section 3.1.2). Spectra were evaluated not only for a nominal incidence but for different source inclination angles. Over the past years, it has been widely acknowledged that the Hopkins assumption—that is, the assumption that diffraction spectra due to different illumination angles can be adequately approximated through translation—is not valid for recent lithographic projection systems (see Section 3.2.6). As shown in Figure 7.32, nine different source incidence angles (sometimes provocatively called “No Hopkins” orders) are considered. The orders are fixed. In the future, a mechanism self-adapting to the optimized source should be devised. In contrast to the former examples, a more recent image simulator was employed (Evanschitzky et al., 2009) that does not employ a fixed source grid but allows for an arbitrary placement of source points.

The optimization study is bulk image-based, using a constant threshold, similarly as before. The threshold is obtained as follows: For each source, generated by the optimizer, the threshold is evaluated by performing a threshold-to-size computation for the dense vertical pitch. The feature of the dense pitch is hence on target for this threshold. In order to account for thin film effects, the CD is extracted at the ten-percent bottom position of the resist on top of an index-matched substrate. A 100-nm thick chemically amplified resist is assumed (with a refractive index of $1.723 - 0.032j$); the nominal defocus position is set to -50 nm (inside of the resist).




Figure 7.32: Nine constant incidence angles are assumed for the computation of the angle-specific spectra. The “No Hopkins” orders are placed such that two orders lie on the center dipole sigma of 0.8.

7.4.2 Optimization Variables

Accounting for the greater flexibility of the alternative image simulator, the polar-coordinate-based source variable definition is slightly adapted, see Figure 7.33. The number of source points is determined




Figure 7.33: Source representation using a polar grid.

by the number of radii and the number of points on the outer radius ($\sigma : 1$). The so-obtained point density is maintained for the inner radii by scaling the outer point number with the radius. The source is defined only in one quadrant exploiting symmetries. For the total number of source points we hence have:

$$n_s = \left(\sum_{i=1}^{n_r-1} \left[\frac{n_p}{i} \right] \right) + 1 \quad (7.71)$$

where n_r is the number of radii
 n_p is the number of outer source points
 $[.]$ denotes the closest integer

Given the number of intensity values n_i , the total number of source combinations is,

$$|\text{sources}| = n_i^{n_s} \quad (7.72)$$

In this paper, we use 20 radii, 20 outer source points and eight intensity values,

$$\begin{aligned} n_p &= 20, n_r = 20, n_i = 8 \\ \rightarrow n_s &= 210 \\ \rightarrow |\text{sources}| &= 8^{210} = 2^{630} \end{aligned}$$

The source can hence, for example, be encoded using ten (64 bit) integers.

This study is concerned with a one-dimensional line/space pattern representation. The OPC measures are thus restricted to size-biasing and scatter bars. Both schemes can be represented parametrically in a straightforward manner. For all seven pitches, separate bias parameters are introduced, each in the range of -10 to 60 nm (mask scale). Because of the binary parameter representation of the GA used here, real variables are not only defined by their range but also by their granularity (step length), which for the size biases was set to 0.4 nm. The SRAF configuration strategy is depicted in Figure 7.34. SRAFs are placed in the area between two features. The maximum width of this area is equal to the pitch subtracted by the line width of the feature. One parameter (α in Figure 7.34) is dedicated to scaling this area between this maximum width and zero. Since SRAFs are assumed to be placed in the center between two features, also the SRAFs area is assumed to be centered. The number of SRAFs to be placed is determined by another variable (n_a). Given the number of scatter bars and the usable area, a (1:1-duty ratio) pitch for the SRAFs (p_a) is determined. The line width of the features is then modified through a scaling parameter (β), also in the range of zero and one. Both scaling parameters have a granularity of 0.1 . No scatter bars were allowed for the dense pitch and the horizontally oriented lines. The maximum number of SRAFs for the remaining pitches is as follows: 125 nm: 2, 135 nm: 3, 270 nm: 4, 315 nm: 5. The total number of size-biasing and scatter bars parameters is 19.




Figure 7.34: The scatter bar configuration for each pitch is determined by three parameters. The scaling of the usable SRAFs area (α), the number of bars (n_a) and the linewidth of each bar relative to the 1:1-duty-ratio pitch (β).

For the projector wavefront control, the coefficients of the according Fringe-Zernike polynomials are varied. Principally, the modification of any combination of (symmetric) Zernike orders is eligible to optimize process window overlaps. Higher-order spherical aberrations, however, can be considered the critical contributors since they are suited to achieve a pitch-dependent focus shift. We have therefore included the spherical aberrations shown depicted in Figure 7.35. As can be seen from the phase maps, the different orders can be expected to show a different effect on the different spatial frequencies arising from different pitches and features sizes. Thus the goal is to find a combination of different orders that maximizes the overlap of the different process windows. Each Zernike coefficient is varied in the range of -300 to 300 m λ ; with a granularity of 1 m λ .

In all optimization examples presented here, the total number of parameters for the source, the mask and the projector is 34.




Figure 7.35: Spherical Zernike orders. Phase maps depicted for 0.5λ

7.4.3 Optimization Objectives

Using a simplistic 1-D definition of the mask, manufacturability criteria are omitted in this example, and only aerial-image-related criteria are considered.




Figure 7.36: Sidelobe criterion. (a) In the case of a sidelobe printing violation, a penalty, proportionate to the intensity distance to a “sidelobe-free” region, is assigned. (b) If no sidelobe printing occurs, the candidate is rewarded according to the distance to an intensity at which sidelobes print.

Like before, criteria are understood as “soft” constraints, where a lower degree of violation enhances the utility of that solution. There may or may not be a turning point at which the solution is actually considered infeasible. These kind of constraints can alternatively be considered as regularization criteria, increasing the uniqueness of solutions of ill-posed problems (see similar penalty-method-based approaches discussed in Section 7.1). A criterion, similar to the band criterion in the former examples, was used (illustrated in Figure 7.36), punishing a candidate solution if its sidelobes bear the danger to print and rewarding solutions that in bright areas show an intensity significantly above the threshold.

In this study, a process-window-centric problem definition is offered. Therefore the maximization of the process window overlap can be considered the central objective. In order to obtain a common process, all windows have to overlap. If only one of the windows does not intersect with any other window, the solution is not feasible. A minimum process window overlap can hence be described as a hard constraint. To allow infeasible candidates to eventually converge toward the feasibility boundary, the distance of those solutions to the process window of the dense pitch, which can thus be thought of as the anchor window, is being minimized (see Figure 7.37(a)). Virtually any well-behaved point in the process window may be suited for the distance comparison. We have used the centroids; an alternative

is the mean of the threshold range of the process window at its best focus. In case all process windows intersect, by whatever small value, the process window criterion becomes a figure of merit that is to be maximized (Figure 7.37(b)).




Figure 7.37: Process window overlap criterion. (a) If any of the process windows do not intersect, the distance of all process windows from the dense process windows is used as a penalty. (b) If all process windows intersect, their common overlap is used as a figure of merit.

The process window overlap constitutes a criterion that incorporates a number of implicit evaluation factors. For example, a large common window also indicates a large exposure latitude and depth-of-focus of the individual processes. Moreover, with the dense pitch serving as an anchor, the CD deviation at the nominal threshold (obtained for the dense case) can be expected to be marginal. Mainly as auxiliary criteria to speed up the convergence in the early GA stage, two additional figures of merit have been included: (1) the explicit minimization of the CD deviation for the different pitches at the nominal threshold and (2) the ILS at the nominal CD. Further criteria such as mask manufacturability constraints could be incorporated in a straightforward manner.

The fitness function of the single-objective replacement problem is

$$\mu : \mathbf{x} \mapsto \sum_{c=1}^4 w_c \mu_c(\mathbf{x}) \quad (7.73)$$

where μ_1 is the sidelobe merit,
 μ_2 is the process window overlap merit,
 μ_3 is the CD deviation penalty and
 μ_4 is the ILS merit
 w_c are the according weight parameters

For the multi-objective formulation, the four criteria μ_1, \dots, μ_4 directly form the sub-objectives:

$$\boldsymbol{\mu}(\mathbf{x}) = [\mu_1(\mathbf{x}), \mu_2(\mathbf{x}), \mu_3(\mathbf{x}), \mu_4(\mathbf{x})]^T, \quad (7.74)$$

The simulation and evaluation flow used to determine the individual objectives is summarized in Figure 7.38.

7.4.4 Results

The optimization behavior of genetic algorithms can be broadly described as an evolutionary process in which individuals are improved by selecting and recombining above-average candidates, eventually

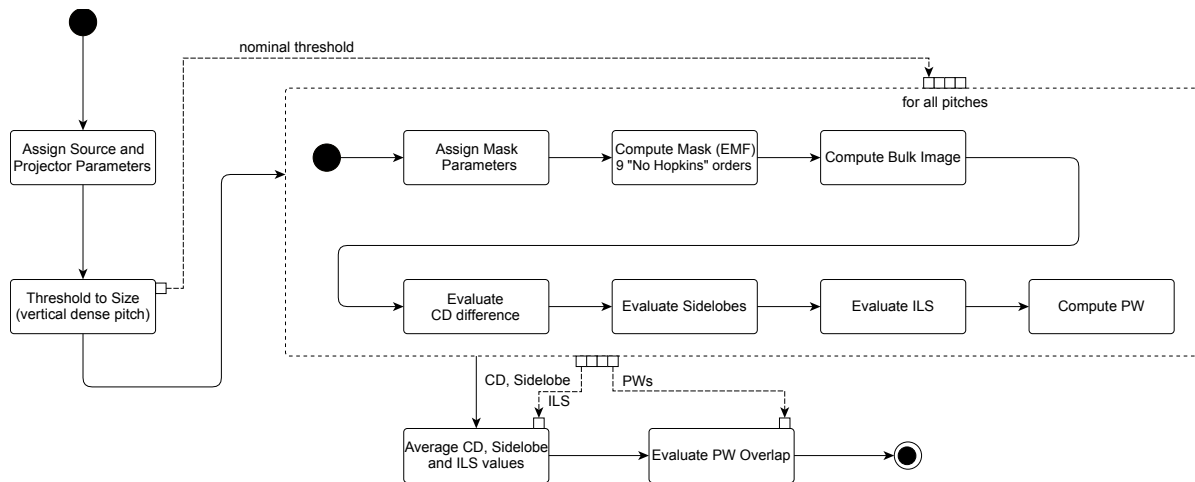


Figure 7.38: Activity diagram illustrating the simulation and evaluation flow for each candidate solution.

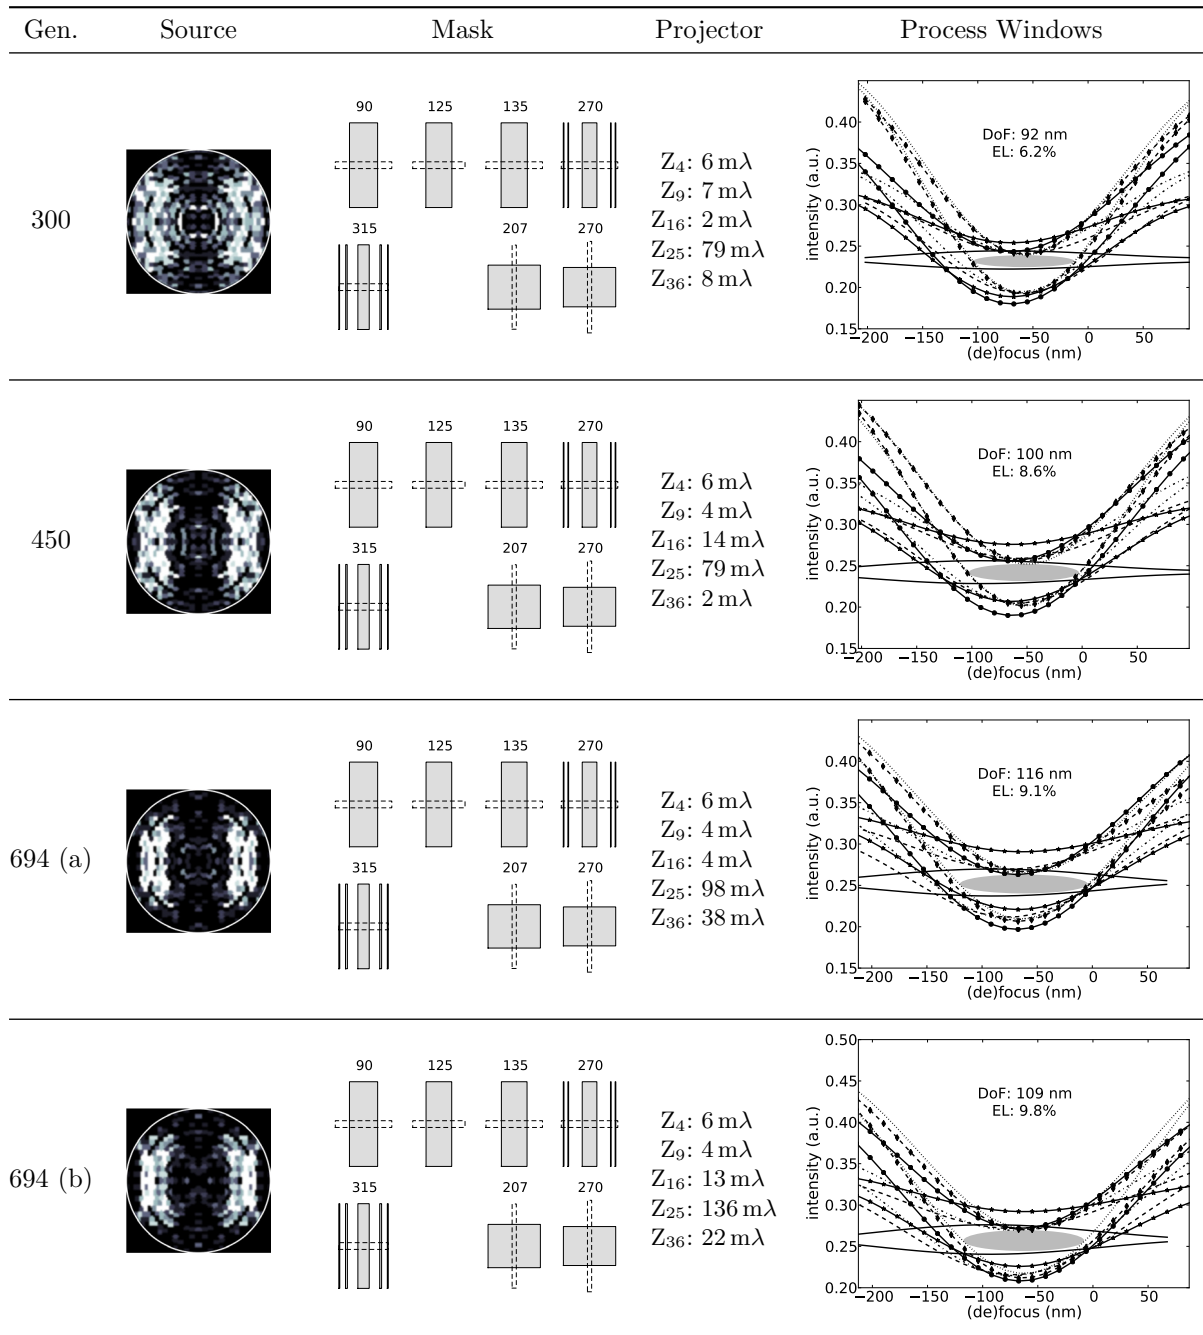
yielding a (local) optimum. Since in addition to the pay-off of the visited search points no other information of the search space or the optimization variables are exploited, the search can be characterized as a probabilistic process, which is illustrated in Table 7.5.

Table 7.5: Evolution of best candidate solutions for a number of example generations. The first column denotes the generation index, the second column shows a Cartesian grid representation of the source, the third column the mask where the numbers are the pitches in nanometers, the fourth column presents the Zernike coefficients, and the last column shows the process windows.

Gen.	Source	Mask	Projector	Process Windows
10		<div> <div>90</div> <div>125</div> <div>135</div> <div>270</div> <div>315</div> <div>207</div> <div>270</div> </div>	Z_4 : 147 m λ Z_9 : 43 m λ Z_{16} : 16 m λ Z_{25} : -46 m λ Z_{36} : -138 m λ	
50		<div> <div>90</div> <div>125</div> <div>135</div> <div>270</div> <div>315</div> <div>207</div> <div>270</div> </div>	Z_4 : 35 m λ Z_9 : 40 m λ Z_{16} : 14 m λ Z_{25} : 80 m λ Z_{36} : 61 m λ	

continued on next page

(continued)



The best result after the tenth generation is displayed in the top row. Both the parameters and the objective values are insignificantly different from randomly generated solutions: The source, for example, does not show any preference for specific frequencies. In this early stage, mainly the distance between process windows is minimized. The large aberration values and SRAFs can thus be attributed to the strong impact they have on this distance. After 50 generations (second row), process windows start to overlap, specific frequencies in the source are accentuated, others begin to diminish. Mainly the sidelobe criterion leads to a SRAF configuration showing a significantly reduced risk of unintentional

printing. In generation 300 (third row), there is already a usable process window. The wavefront control parameters become an effective actuating variable in increasing the common overlap; a tendency of large Z_{25} parameters and significantly smaller values of the other Zernike coefficients emerges. After another 150 generations, further process window asymmetries are resolved, and additionally the exposure and focus latitudes of the individual pitches are increased. This is achieved mainly by an increase of the 16th Zernike coefficient and a pronounced change of the source, strongly emphasizing specific off-axis frequencies. The resulted process window shows both an acceptable DOF and EL. Since in this first optimization study we have employed a weighted sum approach with a niching technique, the final population consists of numerous local optima. We have limited the number of generations to 695 (runtime of about 2 days using 60 processors). The best two final solutions are depicted in the last two rows of Table 7.5. Both solutions show a large overlapping process window both in terms of the usable exposure latitude and the depth-of-focus. The source shapes of both results are similar. As could be expected, both show a pronounced dipole component. The wavefront is modified by a considerable amount of aberrations, especially the 25th and the 36th order Zernikes.

To verify the performance of the final solution candidates, a reference solution with an “aggressive” y -polarized dipole and a simplistic size-biasing OPC was computed. The center sigma for the dipole is 0.8, with an opening angle of 40 deg. In comparison to the first final solution Figure 7.39(a), this “hand-optimized” solution Figure 7.39(b) shows a significantly narrower depth-of-focus, due not only to the principally different through-focus behavior of the different pitches but also to the process window asymmetries caused by the different mask phase responses per pitch. Applying the same aberration parameters resulted from the optimization to the dipole case (c) does not improve performance but instead attenuates process window asymmetries, further reducing the usable window. Removing wavefront control from the final optimization solution (d), on the other hand, introduces a pronounced shift of the best focus and hence leads to a drastically smaller window. Last, replacing the co-optimized source with the standard dipole (e) leads to a complete loss of any process window overlap.

We have conducted the same optimization study using a multi-objective GA. As explained before, this approach reduces the risk of introducing a weight-induced unintentional bias. Broadly speaking, this is achieved by separating the actual optimization procedure from decision making, the final selection of one or more preferred solutions. The latter can be either performed automatically, in a post-optimization process, or, as in this work, through manual user selection. This allows for an exploration of different, equally well performing solutions. Table 7.6 presents three examples for final solutions after a 1000-generation run, which had a runtime of about 3 days (using 60 processors). Neither of the best solutions in this (nor in the previous) optimization run suffered from sidelobe printing. The first solution (a), with a large common DOF and EL, exhibit large values of the Zernike coefficients 16, 25 and 36. The source shows some similarity with the former results obtained with the single-objective GA, albeit with a reduced opening angle of the dipole-like structure. This increases the overall threshold intensity and the exposure latitude of the dense pitch. The additional inner, narrower dipole can be speculated to increase the exposure latitude of the larger pitches. The solution depicted in the second row (b), in contrast, displays a slightly less pronounced removal of window asymmetries. The wavefront is modified through a larger number of slightly smaller Zernike numbers. This renders the practical feasibility of this solution more likely than candidates with extreme Z_{25} coefficients. In addition, the source shows a more pronounced inner dipole and poles on the y -axis, significantly increasing the process window of the horizontally oriented lines. A similar behavior can be observed in the last example (c). There, however, significantly lower aberration values are obtained, leading to a less tight overlap of the process windows in comparison with the former solutions. The usable common process window is still not prohibitive, indicating that source/mask/projector pupil optimization (SMPO) does not necessarily have to translate all process windows into the exact same focus regime.

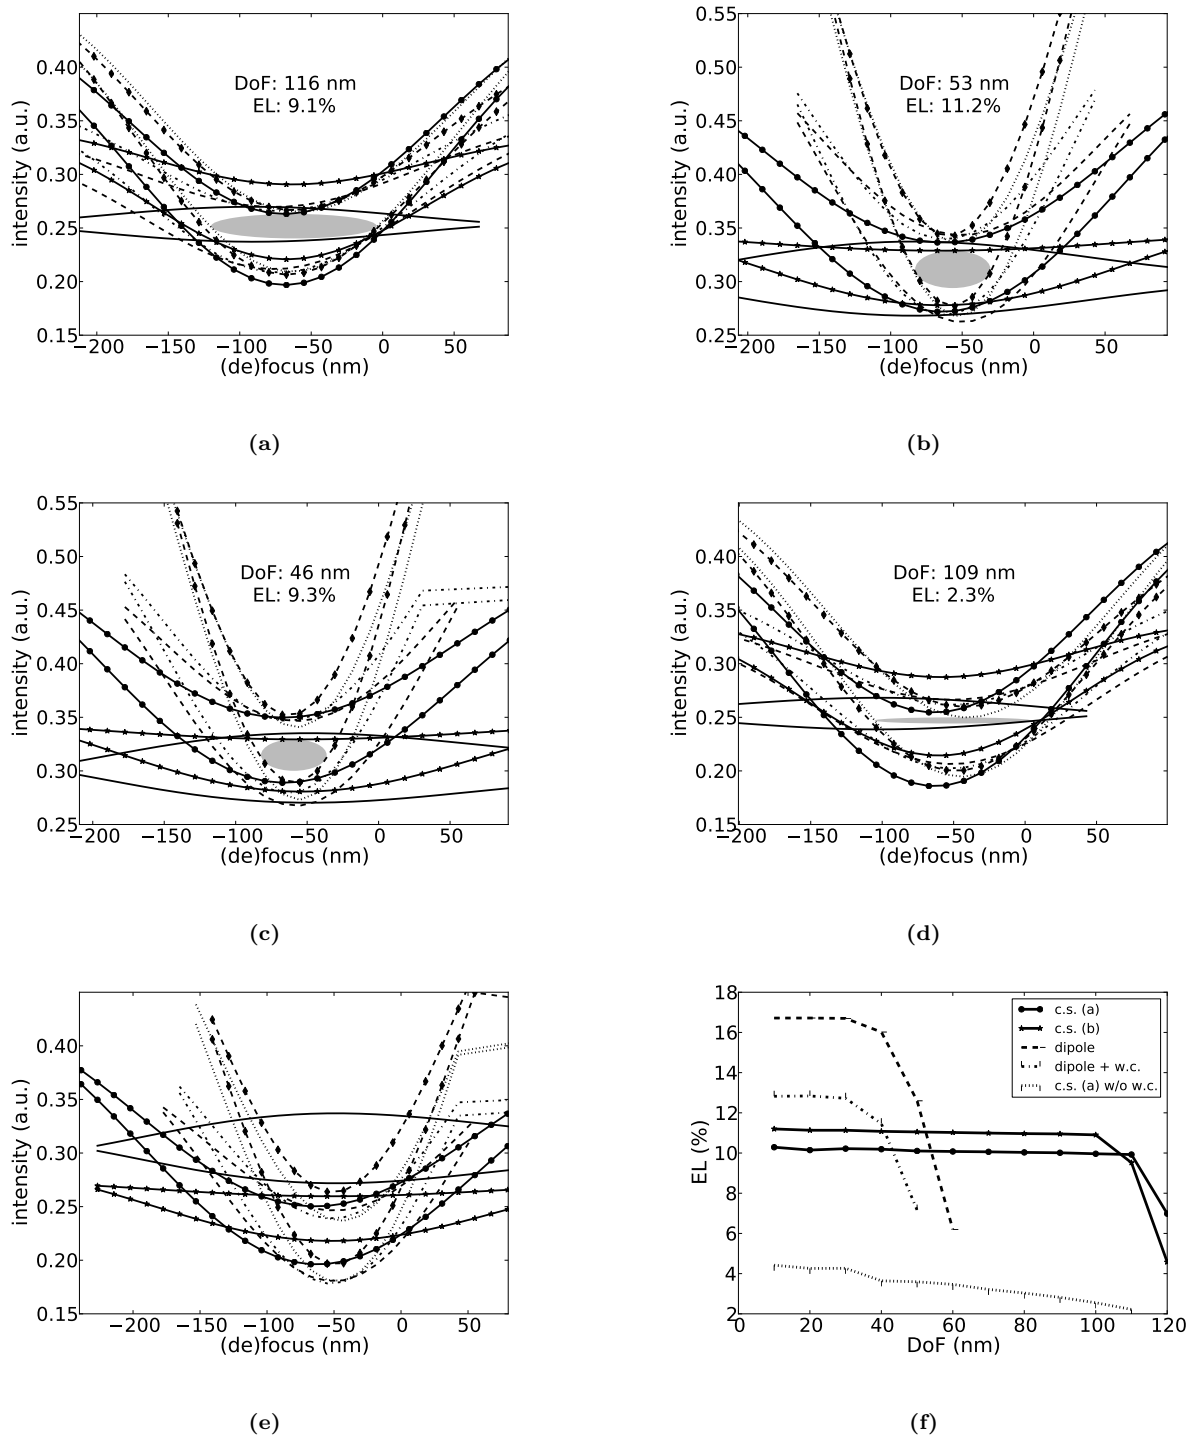
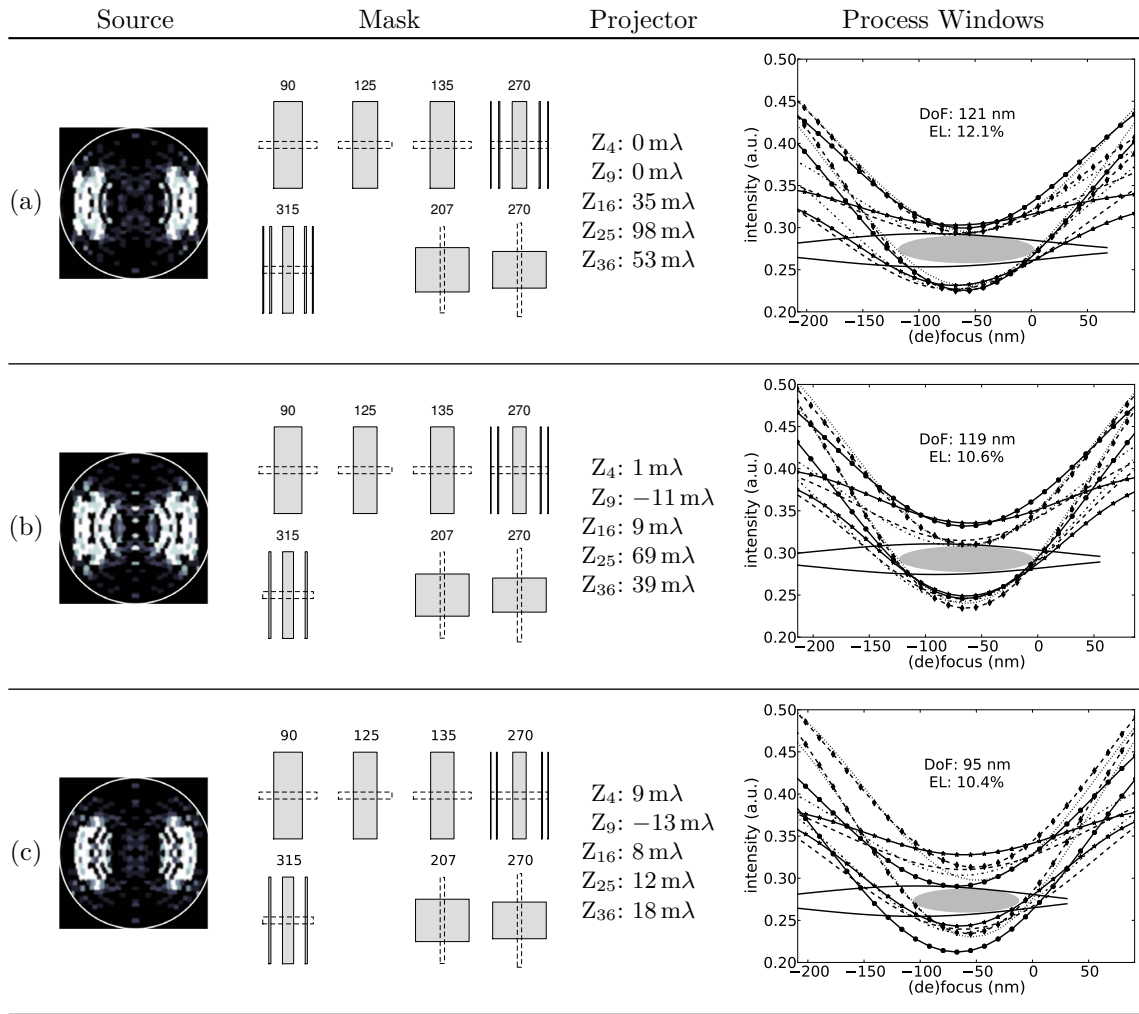


Figure 7.39: Interdependence of the source, the mask and the wavefront configuration. (a) Final optimized solution; (b) “Hand-optimized” dipole and size-biased mask solution; (c) “Hand-optimized” option incorporating the wavefront control parameters of the rigorously optimized solution; (d) Final optimization result without wavefront control; (e) Final optimization result with wavefront control using the “hand-optimized” dipole instead of the optimized source; (f) Depth-of-focus versus exposure latitude performance for the different options [c.s. (a) and c.s. (b) denote the alternative final candidate solutions]. For a legend for the process windows, see Table 7.5.

Table 7.6: Example solutions of the multi-objective genetic algorithm. For a legend for the process windows, see Table 7.5.

8 Mask and Wafer Stack Optimization

The performance of the lithographic process is of course determined by a much larger number of parameters and conditions than those optimized for in the last chapter. For example, the wafer stack has been accounted for only superficially by making simple assumptions about its thin-film layers and by regarding the photoresist as a simple threshold detector. Although we have also demonstrated the extensibility of the proposed direct optimization approach, such as the incorporation of the projector pupil domain, an attempt to optimize virtually all aspects of the lithographic system is neither possible nor useful. Specifically, the feasibility of such an approach is restricted by its computational cost, which can be expected to be extremely high; but also its utility is limited since, because of the multitude of influencing and interacting factors, results are extremely hard to interpret, drastically impeding the extraction of suited conclusions. In this chapter we thus restrict our perspective to specific parts of the lithographic system. In the first three sections, the main emphasis is put on the absorber stack of the mask, both in optical and EUV lithography, the last section is aimed at an optimization of the wafer stack in terms of its anti-reflective performance under multiple exposure conditions.

8.1 Mask Absorber Optimization for Line/Space Patterns in Optical Lithography

In the last section of the previous chapter, we have demonstrated the strong impact of EMF effects on the lithographic process behavior. This can be mainly attributed to the topography-induced phase effects, which lead to a significant displacement of process windows in comparison with a Kirchhoff mask approximation. We have shown how projector aberrations can be utilized as a compensation measure. Of course, the problem can be also attacked by adapting the origin of the problem, the mask. In this section, we therefore present an optimization study aimed at optimizing the mask stack in order to improve the exploitable process range and especially to maximize the common process window of different features, which can be considered a difficult task since the mask-induced phase variations differently affect individual diffraction orders. In this example—summarizing the study of Erdmann et al. (2007)—we therefore again concentrate on a one-dimensional case, in which a dense and an isolated line/space pattern are to be simultaneously imaged. The target feature is a dark line with a width of 40 nm. The pitch of the isolated case is 400 nm, and a 1:1 duty ratio is used for the dense case. Traditionally, binary or AttPSMs consist of single absorber layer such as molybdenum-silicon (MoSi). Comprehensive studies on alternative materials have been conducted, for example, by Kitahata et al. (2006).¹ By introducing an additional layer, the absorber can be expected to exhibit a pronouncedly higher flexibility in terms of adaptations aimed at compensating for phase or polarization deficiencies (Yoshizawa et al., 2006). The following optimization experiment hence includes bi-layer configurations.

8.1.1 Simulation Conditions and Optimization Variables

The optimization variables are depicted in Figure 8.1. They include the thickness of the two absorber layers (D_1 , D_2), their material properties and the widths of the main feature (w), the distance between main and assist feature (d_a) and the width of the assist. The following constraints are applied:

$$w_a, d_a \geq 10 \text{ nm.}$$

¹ See also Section 2.6.3 for a discussion.

As usual, assist features are placed in order to reduce the iso/dense bias and are hence only introduced for the isolated case. The total thickness of the absorber is limited to 200 nm, thus:

$$D_1 + D_2 \leq 200 \text{ nm.}$$

To enforce the feasibility of solutions, the following simple decoder strategy is employed. An auxiliary optimization variable D denoting the total fitness is introduced, with $0 \leq D \leq 200$ nm. The thickness values of the individual layers are then expressed as a ratio γ . Specifically,

$$\begin{aligned} D_1 &= \gamma D \\ D_2 &= (1 - \gamma)D \quad \gamma \in [0, 1]. \end{aligned}$$




Figure 8.1: One-dimensional bi-layer absorber stack optimization. The optimization variables include the material properties (M_1, M_2) of the absorber layers and the individual thickness values (D_1, D_2). Aiming at the co-optimization of a dense (80 nm) and an isolated pitch (400 nm), an assist feature placement—parametrically determined its width (w_a) and the distance from the main feature—is additionally incorporated for the isolated case, in which additionally a size biasing is performed (parameter w).

In this first example, the material property parameters of the absorber layers (M_1, M_2) are specified through the (n/k) (refractive index/extinction) tuple, the range of which is set to:

$$0.2 \leq n_i \leq 3.0 \quad \text{and} \quad 0.0 \leq k_i \leq 3.0 \quad (i = 1, 2).$$

In addition to the standard i193 setting (see Section 7.4), two fixed illumination conditions are employed (see Section 2.6.2): A fixed CQuad with an aperture angle of 20 deg. and a dipole with an aperture of 40 deg., both with inner and outer radii of 0.8 and 0.96, respectively, both x/y-polarized. To account for thin-film effects, the same resist material settings as in the previous example are employed. Like there, the index of the substrate is matched against the resist, thus emulating a perfect BARC. A constant threshold resist model is applied. The EMF simulations were again conducted using the waveguide method with incidence-angle dependent spectra, centered about the poles constituted by the illumination source.

8.1.2 Optimization Criteria

The first category of optimization criteria aims at the imaging performance due to the optimized mask. Thus similar, yet subtly different, criteria as in the previous examples and additional mask-related criteria are employed.

- (1) The first objective is the maximization of the *global contrast* (see (2.12)) of the image of the dense pattern.
- (2) Exposure latitude is maximized at a fixed DOF of 200 nm in the dense and 50 nm in the dense/isolated co-optimization case. The small DOF in the latter case pays tribute to the fact that in this study the iso/dense bias is accounted for by only a single assist feature, which is prohibitive for a wider DOF.
- (3) A sidelobe criterion like in the former cases is introduced: Here, the threshold maximizing DOF (best threshold) is obtained from the process window, increased by thirty percent and applied to the aerial image. Sidelobes are assumed to print if they fall below this threshold. Sidelobe exhibiting candidates are penalized proportionately to the intensity difference between the sidelobe and the former threshold.
- (4) MEEF (see Section 2.2.5) is accounted for in simplistic manner, by evaluating the three process windows: at the nominal mask size and for the mask with a size respectively decreased and increased by 1 nm. The overlap of the resulting PWs is to be maximized.
- (5) A similar strategy is used in order to minimize the sensitivity of the mask against polarization variations. There, the intensity in the preferred polarization state (IPS) metric is used (Geh et al., 2005), which in the y-polarization case is defined as follows:

$$\text{IPS}_y := \frac{s_0 - s_1}{2s_0}.$$

where s_0 and s_1 are the respective Jones parameters (see Section 3.2.3). An IPS of 0 means that none, a value of 1 that all of the light is correctly polarized. Two process windows, for a value of 1 and 0.9, are evaluated and their overlap is to be maximized for this criteria.

- (6) The definition of materials through a list of refractive indices and extinction coefficients renders the approach highly flexible but also bears the danger of yielding material combinations that are not available or feasible. To circumvent this issue, we have devised a technique assigning all materials in the n/k landscape a suitability factor. First, a list of materials known to be feasible is provided, obtained in this case from RIT Nanolithography Research Labs. A materials feasibility is then given as:

$$\mu_M(n, k) := \sum_i \min(p_i(n, k), 1), \quad (8.1)$$

where p_i denotes the probability that the i th element of the provided materials can be modified such that it resembles the properties n, k .

$$p_i(n, k) := \exp \left[-\frac{1}{2} \left(\frac{(n - n_i)^2}{\Delta n_i^2} + \frac{(k - k_i)^2}{\Delta k_i^2} \right) \right],$$

where $\Delta n_i, \Delta k_i$ can be considered elasticity values, determining how well a material can be modified. Here, for simplicity, we assume a constant value of 0.1 for both components of all materials. Figure 8.2 shows the map of materials considered in this study and the corresponding feasibility distribution due to (8.1)

- (7) Also the last criterion is related to mask. It is concerned with its inspectability, which is drastically degraded under an increased reflectivity of the mask blank. A simple penalty criterion is used here, punishing candidates with a reflectivity of above 20 percent (at 193 nm).

8.1.3 Results

The first results are obtained for the dense line/space layout and the dipole only. The contrast, exposure latitude and the reflectivity criteria are applied using a single-objective GA and a weighted sum

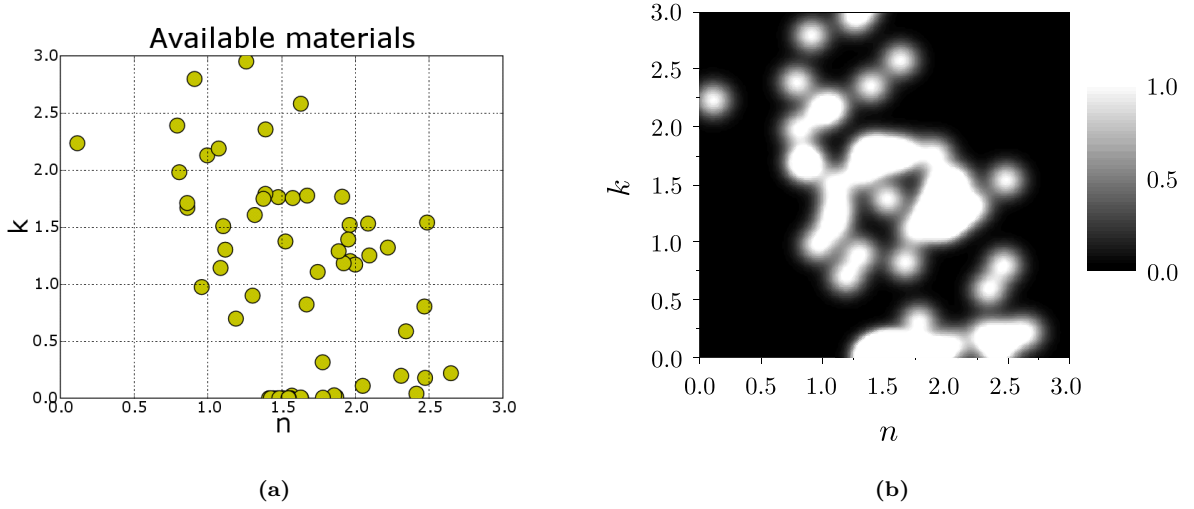


Figure 8.2: Absorber material availability landscape: (a) Raw material data (taken from RIT Nanolithography Research Labs); (b) Feasibility map according to (8.1).

replacement problem. The number of individuals is set to 60, and the maximum number of generations is 1000. Both restricted tournament selection and Gray code are employed. Two-point crossover with a probability of 0.6 and a mutation probability of 0.005 are used. As a reference, the absorber stack presented by Yoshizawa et al. (2006) is simulated and shown in Figure 8.3(a) and (b). The threshold is normalized such that the best threshold (which maximizes the DOF) corresponds to one. The two-beam imaging condition presented by the combination of a dipole and the dense line/space configuration principally provides a large DOF, provided the zeroth and the first orders are balanced accordingly (see the respective example in Section 2.6.3). Evaluating the process window not only for the nominal line width but also for a mask bias of ± 1 nm (wafer scale) reveals a slightly lower sensitivity of the Ta/SiO₂ bi-layer against linewidth variations, when compared to the performance of the standard MoSi absorber. This minor difference, however, amounts to a considerable increase of the usable process window. The optimized solutions (Figure 8.3(c) and (d)) show a similarly large DOF, apparently maintaining a similarly well-behaved balance of diffraction orders. Not optimized for the MEEF criterion, however, the left (c) candidate exhibits a notably aggravated sensitivity against CD non-linearities, demonstrated by the process windows for the ± 1 nm bias. Including the MEEF criterion (d), compensates for this, albeit at the price of a drastically reduced threshold and consequently of an incompatibility with less dense structures.

A similar study, albeit using the CQuad instead of the dipole, is conducted with the memetic algorithm with function approximation (cf. Popp, 2007),² using the same GA configuration as before. In addition, a neural network with seven input nodes for the optimization variables and a hidden layer with nine nodes and a single output node (fitness function value) is employed.³ The iRProp training algorithm is used (Igel and Hüsken, 2000). As a local optimizer the Nelder-Mead algorithm is applied. Using the original mapping of parameters to the ANN input (see (5.29)) led to fast convergence during learning but a disappointing convergence behavior. As pointed out by Popp (2007), this is due to the fact that for a large number of well performing candidates the upper bound constraint for the total absorber thickness is active (cf. Figure 8.5 of the next example, also using a CQuad source). As a consequence, the ANN representation of the fitness function is dramatically flattened, leading to a large number of prematurely terminated local optimization attempts. To alleviate the problem, a scaling accounting

²For a discussion, see Section 5.6.

³The configuration of the hidden layer is validated through a number of numerical experiments in which the neural network is trained with off-line data obtained from GA test runs (Popp, 2007). The same test is also used for the determination of an appropriate training algorithm.

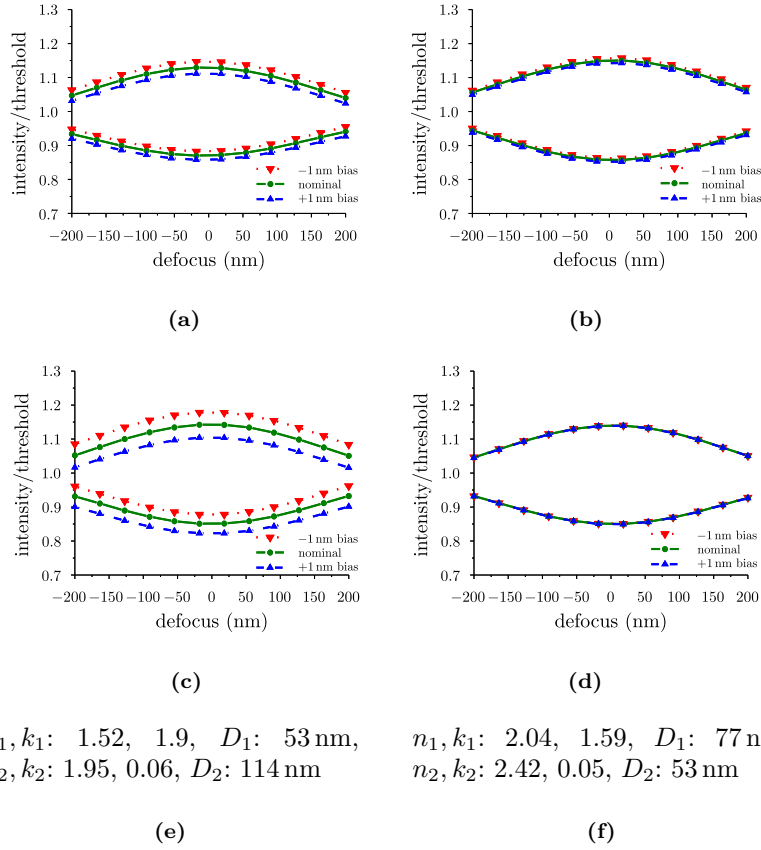


Figure 8.3: Results of dense line/space pattern absorber stack optimization, showing process windows for the nominal mask size (circles), 1 nm smaller lines (wafer scale; flipped triangles) and 1 nm larger lines (wafer scale; up-right triangles) on the mask. The solutions (a) and (b) are obtained with the standard MoSi and the Ta/SiO₂ (Yoshizawa et al., 2006) absorber stack, where the latter shows a less severe sensitivity against line width variations. (c) Solution optimized for exposure latitude for the nominal mask line width, exhibiting a pronounced sensitivity against line width variations, which can be drastically reduced by incorporating the MEEF criterion into the optimization (d). The settings corresponding to (c) and (d) are shown in (e) and (f), respectively, presenting the material properties and the thickness values of the absorber stack $[n_i, k_i, D_i \ (i = 1, 2)]$.

for the frequency of parameter values in the current population is devised: Given the mean parameter value \hat{x} of the current population and its standard deviation σ_x , then we have for the updated scaling mapping

$$[\hat{x} - 2\sigma_x, \hat{x} + 2\sigma_x] \rightarrow [-1, 1].$$

With this adaptation, the MA with function approximation showed a slightly improved convergence rate compared with the GA. This is depicted in Figure 8.4. The graph additionally compares the performance of the non-approximated MA. In both cases, a constant trigger, with an interval of 15 generations and an egalitarian selection with a probability of 0.1 are adopted (see Chapter 5 for a discussion). The MA with function approximations presents an insignificant computation time overhead compared with GA. Specifically, it requires the training of the network, the local search—where each function evaluation is extremely efficient in comparison to one in which a lithography simulation step is conducted—and the function evaluations due to the model-control subsequent to each local search. In contrast, the non-approximated MA conducts its local search on the original search space. With an average of about 100 function evaluations per local search, each generation at which the local search is triggered corresponds to about 10 GA generations in terms of function evaluations.




Figure 8.4: Convergence of comparison of memetic algorithm with and without approximation (dashed and dotted lines, respectively) and genetic algorithm (solid line) for absorber stack optimization of dense line/space pattern. The mean performance of five runs is shown by the thick lines, the light-shaded regions depict the best and worst performance.

The five independent runs conducted with the function-approximated MA are tabulated in Table 8.1.

The next simulation study is aimed at minimizing the sensitivity of the dense feature against polarization variations, thus employing the IPS criterion. An IPS different from one may for example occur under illumination conditions constituting a compromise for different features and polarization conditions deviating from the nominal x/y-polarization setup. To model such a situation, the CQuad illuminator configured as explained before is used in this example. The optimization results are depicted in Figure 8.5. As a reference solution, the MoSi stack is shown on the left (a), performing relatively well both in terms of exposure latitude and low polarization sensitivity. The optimized solutions, however, show a considerably larger exposure latitude. Both candidates are champion solutions of different optimization runs: The center solution (b) is obtained accounting for the material availability constraint, which is omitted for the solution on the right of the figure (c). There, the pronounced increase in EL

Table 8.1: Dense line/space AttPSM stack optimization using a memetic algorithm with function approximation.

n_1	k_1	n_2	k_2	w (nm)	D_1 (nm)	D_2 (nm)	contrast (percent)	exposure latitude (percent)
1.5	-0.23	1.82	-0.09	52	130	70	81.14	24.23
0.87	-0.94	1.69	-0.18	53	11.22	175.78	80.99	24.18
1.56	-0.28	1.83	-0.19	55	152	48	79.29	23.67
0.86	-0.99	2	-0.05	47	13.42	108.58	81.17	24.25
1.40	-2.30	1.95	-0.11	48	2.36	115.64	80.04	23.91

is achieved through a dramatically smaller refractive index of the absorber top layer, associated with an unbalanced diffraction efficiency for the x-polarization state. The non-availability of an according material renders this solution infeasible, however.

The last example combines the dense and the isolated line/space features. To counteract the iso/dense bias, the optimization includes a variable main feature width and the placement and width variation of the assist feature. The aim is to maximize the common process window and the contrast of the dense feature, while at the same time maintaining a reflectivity of the blank below 20 percent and preventing sidelobe printing. Figure 8.6 presents the aerial images, the common process window and the material properties of the (OPCed) reference MoSi stack, a solution subjected to the material availability criterion and a final result of an only material-range-restricted optimization. The results confirm the former observations of a significant increase of the usable process window when applying an optimized instead of the standard (MoSi) stack. Again, the candidate with an unconstrained material selection considerably outperforms results obtained from a material-availability-aware optimization run. Like in the example before, a thin absorber top layer with a small index of refraction and a moderate extinction coefficient yields the best performance. Also note the differences in the main and assist feature configurations, which can be attributed to a response to the different diffraction order efficiencies.

8.2 Mask Absorber Optimization for Contact Holes in Optical Lithography

A similar problem formulation is also applied to the optimization of contact holes using AttPSMs (the results of which were first published by Erdmann et al., 2008). In the following example, optimum absorbers for non-quadratic holes, with a dimension of $65 \text{ nm} \cdot 90 \text{ nm}$, which can be considered a critical feature for lithography, are sought. A dense (duty ratio 1:1) and a semi-dense setup ($260 \text{ nm} \cdot 360 \text{ nm}$) is co-optimized for. The layout is depicted in Figure 8.7.

8.2.1 Optimization Parameters

A parametric source optimization is conducted, using the same illumination set-ups as for the respective representation in the previous example (see Section 7.2.1).

8.2.2 Simulation Conditions and Optimization Variables

The mask layout parameters are summarized in Figure 8.7. The topology parameters consist of the lateral dimensions (s_x, s_y) of the openings, the thickness D and the material properties, for which three selection modes have been devised:

- (1) *List of user-supplied stacks.* The optimization variable is an integer index l_s into a lookup table. In this study, four pre-defined stacks are considered: Two standard stacks (MoSi and NTAR7) and two additional two-layer stacks proposed by Yoshizawa et al. (2006): Ta/SiO₂ and Cr/SiON.

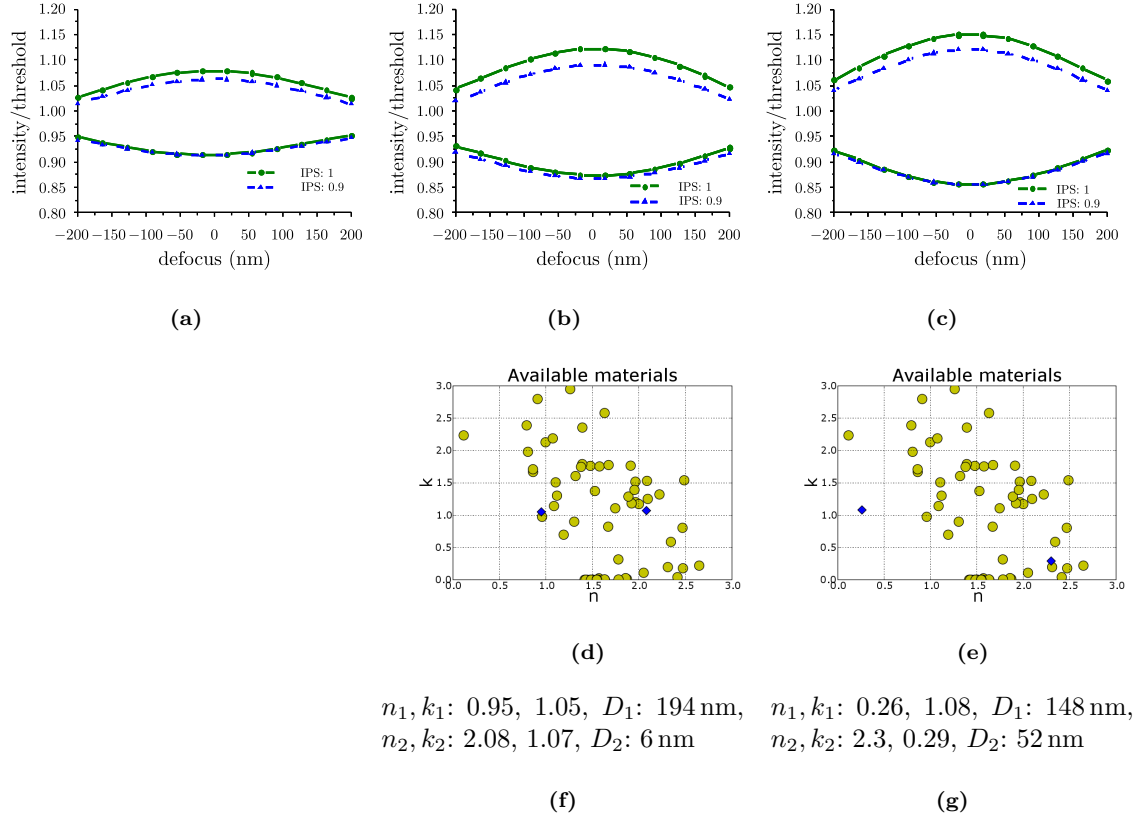


Figure 8.5: Results of dense line/space pattern optimization to minimize polarization sensitivity. Upper row (a)–(c) shows the common process window of solutions evaluated at an IPS of 1 (circles) and an IPS of 0.9 (dashed line). In the first column the MoSi standard stack is shown for reference. The second column (b) depicts the best result of an optimization conducted with materials constrained according to feasibility map as show in Figure 8.2. The solution in the last column (c) is obtained under the unconstrained material mode. The materials corresponding to (b) and (c) are shown in the center row (d),(e). Yellow circles indicate the materials from the database, blue diamonds show the setting resulted from optimization. The according settings, the material properties and the thickness values of the absorber stack $[n_i, k_i, D_i \ (i = 1, 2)]$ are tabulated in the last row (f), (g).

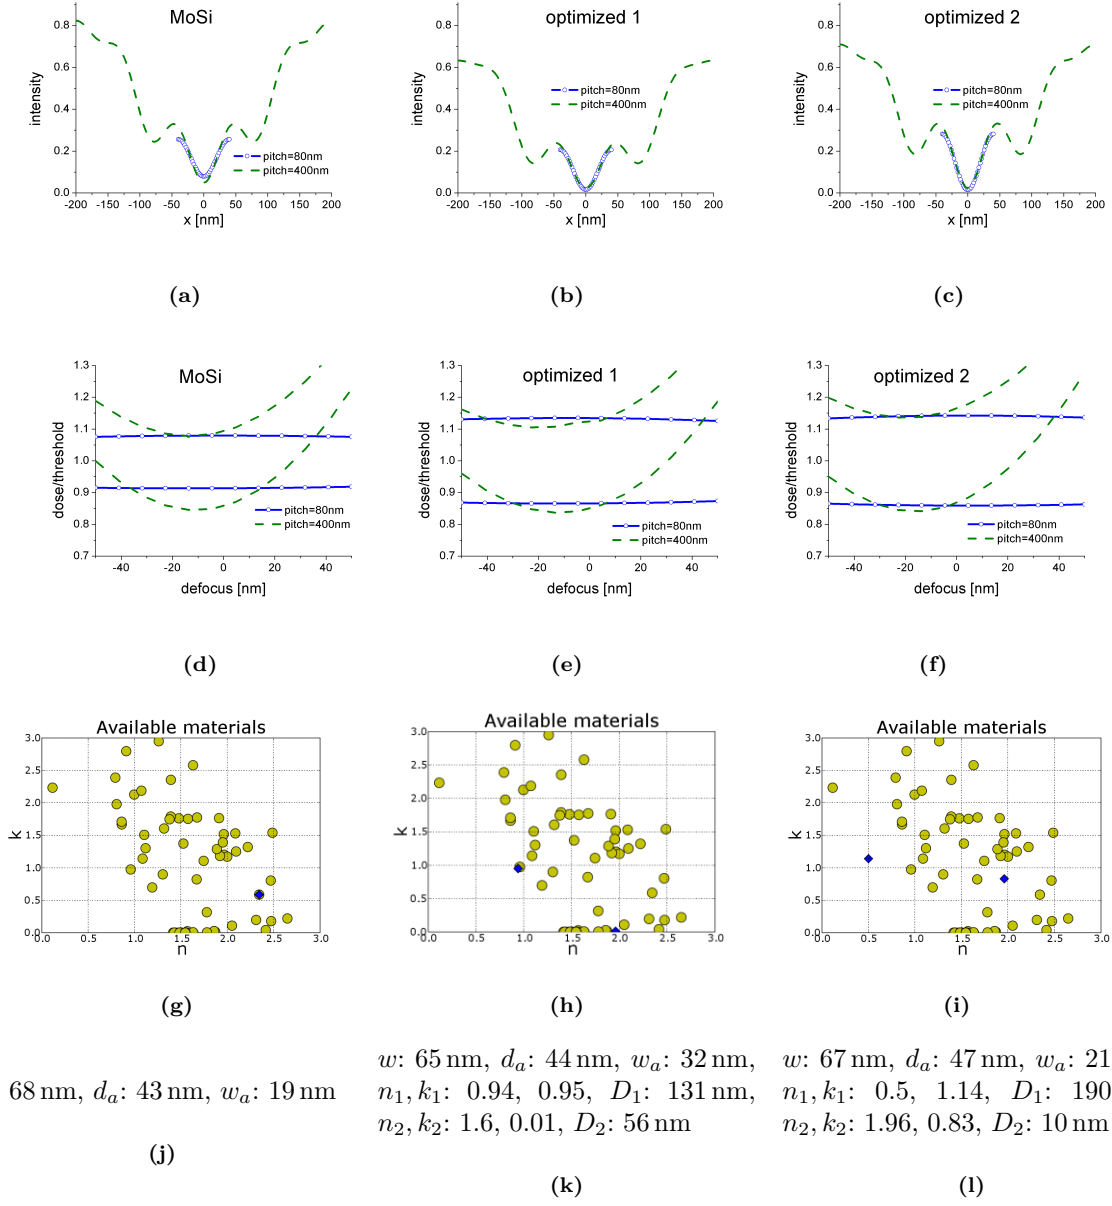


Figure 8.6: Results of dense/isolated line/space pattern co-optimization using different material modes. The first row (a)–(c) shows a crosscut of the aerial images of the dense (circles) and the isolated lines (dashed line). The second row (d)–(f) demonstrates the common process window. The optimized materials are shown in the third row (g)–(i); yellow circles reference denote materials from the data base, the material of the candidate are indicated by blue diamonds. The resulted parameters are summarized in the bottom row (j)–(l), where the material properties and the thickness of the first and the second layer of the absorber stack are n_i, k_i, D_i ($i = 1, 2$), and the parameters w, d_a and w_a denote the width of the main feature, the assist distance and the assist width in the isolated case, respectively. The results in the first column are for the standard MoSi stack, the center column demonstrates the performance of the result obtained with the constrained materials and the right column shows the best solution for the unconstrained material mode.

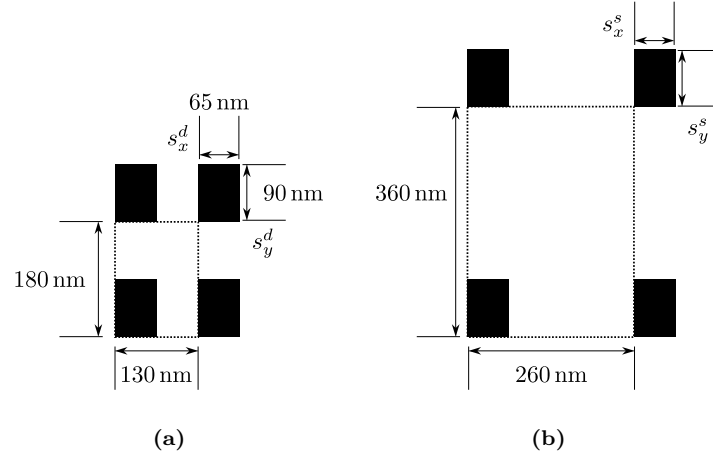


Figure 8.7: Configuration of the non-quadratic (a) dense and (b) the semi-dense contact array. The mask line width are subject to optimization. Specifically, the dense (s_x^d, s_y^d) and the semi-dense (s_x^s, s_y^s) dimensions can be varied individually to compensate for the iso/dense bias.

- (2) *Material data base.* In this mode, the stack is assumed to be a bi-layer stack. For that purpose, the maximum thickness of the absorber is constrained to a maximum of 150 nm, using the same decoder technique as before, and the thickness of two the individual layers D_1, D_2 are subject to optimization. In order to guarantee the availability of materials, an according user-provided data base is employed (in this example, again using the data from RIT Nanolithography Research Labs), and the material of each of the absorber layers is represented by an index into this data base l_{m_1}, l_{m_2} .
- (3) *Free material properties.* This mode is equivalent to the technique in the previous example, yielding four material n_i, k_i ($i = 1, 2$) and two thickness variables D_1, D_2 , where again the n/k ranges are restricted according to

$$0.2 \leq n_i \leq 3.0 \quad \text{and} \quad 0.0 \leq k_i \leq 3.0 \quad (i = 1, 2),$$

and the total absorber thickness is limited to 150 nm.

In addition to the mask topology-related optimization parameters, a simplistic, parametric source co-optimization is employed, adding four additional parameters to the problem (cf. Section 2.6.2):

- (1) the type of illumination t , which is one of annular, CQuad and Quadropole,
- (2) the inner σ_i and
- (3) outer radius σ_o of the source,
- (4) the opening angle α of each pole.

Giving additionally rise to the following feasibility constraints:

$$\begin{aligned} \alpha &\leq 45 \text{ deg.}, \\ \sigma_o &\leq 1, \\ \sigma_o - \sigma_i &\geq 0.2. \end{aligned}$$

which are enforced through the box constraints and a decoder technique identical to the absorber thickness constraint. An x/y-polarization state is assumed.

The same simulation conditions as in the previous example are employed. Particularly, all mask simulations are of course conducted through the rigorous waveguide EMF method. “No-Hopkins” orders are placed such to account for all major incidence angles due to the source.

8.2.3 Optimization Criteria

The goal of this optimization study is the maximization of the common process window of all features, accounting for both the x and the y direction. Thus a process window overlap criterion as explained in Section 7.4 is employed for the following four CD cut-lines:

- dense feature: x and y direction,
- semi-dense feature: x and y direction,

8.2.4 Results

In a first numerical experiment, a co-optimization of a dense and a semi-dense contact hole array using an AttPSM and the material list option is conducted. The optimization parameters thus include the two material indices, one per layer, the two layer thickness values, the four illumination parameters and four feature dimensions, the width and height of the dense (s_x^d, s_y^d) and the semi-dense case (s_x^s, s_y^s), respectively. A population of 90 individuals and a maximum generation number of 1500 are employed. As before, the single-objective GA is configured to use restricted tournament selection, Gray code, a two-point crossover with a probability of 0.6 and a single point mutation with a rate of 0.005. The convergence of a single run is shown in Figure 8.8. There, the fitness is proportionate to the common exposure latitude of the common process. The GA very rapidly produces candidate solutions exhibiting


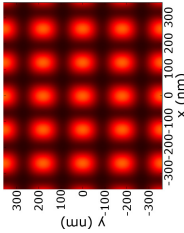
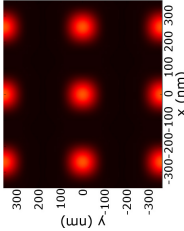
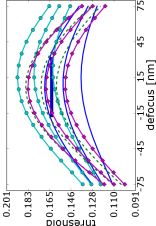
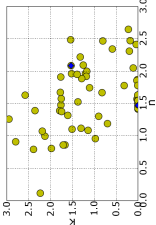
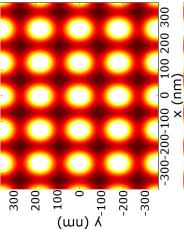
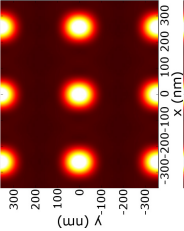
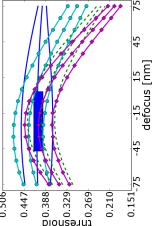
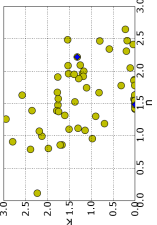
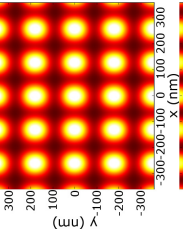
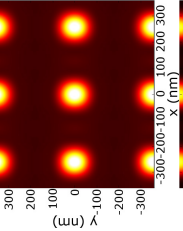
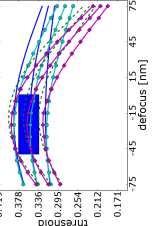
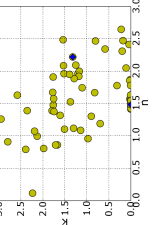
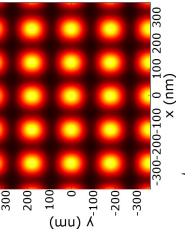
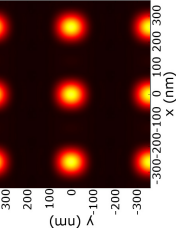
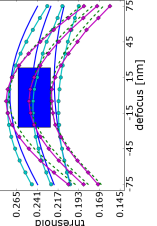
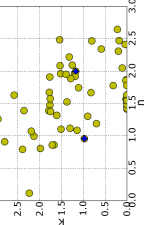


Figure 8.8: Convergence graph of dense/semi-dense contact hole co-optimization. The graphs show the progress of the best (solid line), the average (dashed line) and the worst (dotted line) individual.

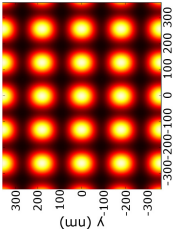
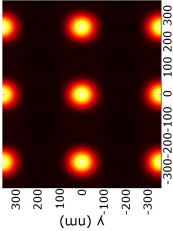
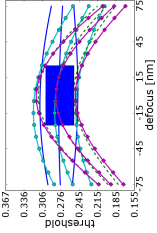
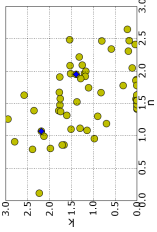
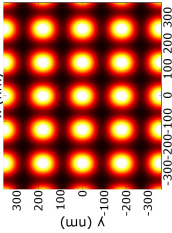
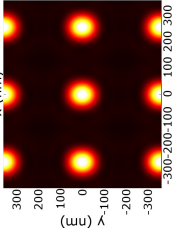
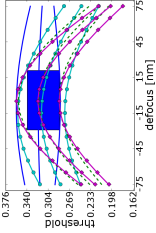
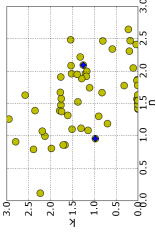
acceptable processes. Table 8.2 illustrates the evolution of best solutions for a number of example generations. After about 50 generations a feasible EL is obtained. Like before, the target DOF is fixed to 50 nm to account for the absence of a sophisticated assist feature placement routine. It is interesting to note that despite the lack of an according criteria, the maximization of process window overlaps appears to reduce individual process window asymmetries.

Table 8.2: Evolution of best candidate solutions for a number of example generations of the co-optimization of dense/semi-dense contact holes. The first column denotes the generation index, and the source settings, the second column shows the aerial image of the dense, the third column the isolated contact hole image. The windows overlap of the two processes is shown in the third column, where a 50-nm DOF is assumed, at which the maximum EL defect is shown. The fourth column shows the materials assigned to the layers (blue diamond) and its location in the material map; available materials are indicated with yellow circles. The last column, shows the EL in percent.

Settings	Dense Image	Semi-dense Image	Process Window	Material	EL (%)
gen.: 0 source: annular $\sigma_i, \sigma_o: 0.81, 1.0$					1.4
gen.: 6 source: quadrupole $\sigma_i, \sigma_o: 0.35, 0.65$					6.3
gen.: 8 source: quadrupole $\sigma_i, \sigma_o: 0.32, 0.7$					12
gen.: 41 source: annular $\sigma_i, \sigma_o: 0.33, 0.79$					15.4

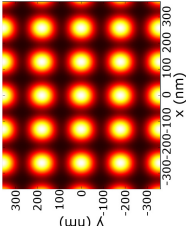
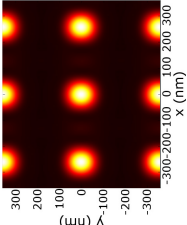
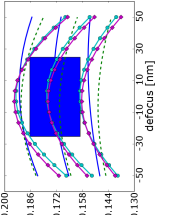
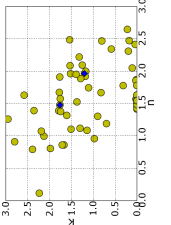
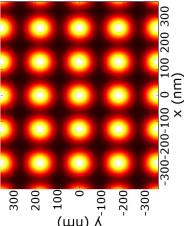
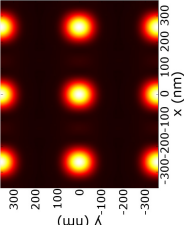
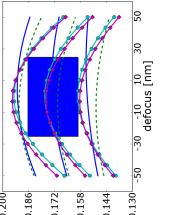
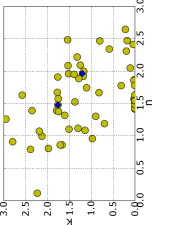
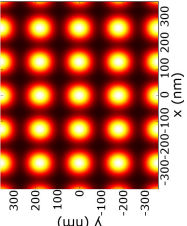
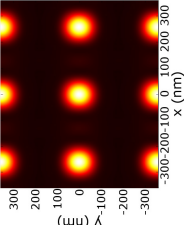
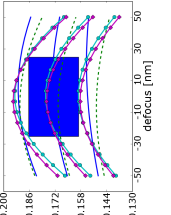
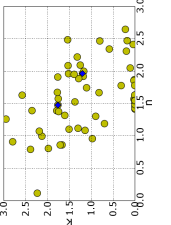
continued on next page

(continued)

Settings	Dense Image	Semi-dense Image	Process Window	Material	EL (%)
gen.: 736 source: quadrupole σ_t, σ_o : 0.38, 0.64					16.7
gen.: 1201 source: quadrupole σ_t, σ_o : 0.4, 0.62					17.2

In order to assess the performance of the optimization result, two comparisons with the remaining alternative material selection modes are performed. The best solutions of each run are shown in Table 8.3. The top row displays the best candidate obtained under a fixed user-provided stack list, the row in the middle demonstrates the best result of Table 8.2, using the material data base, and in the last row the final solution of an optimization run conducted with the free n/k range is depicted.

Table 8.3: Comparison of the best individuals of dense/semi-dense contact hole co-optimizations using different material modes. Like in Table 8.2, the columns show the settings, the dense, the semi-dense images, the common process window and the exposure latitude. The parameters in the first column consist of the absorber thickness D , the illumination type and the illumination variables. In the case of a quadrupole source, α denotes the aperture angle of the poles. The results are obtained with the fixed stack selection mode (top row), the material data base mode (second row) and the free material property mode (bottom row).

Settings	Dense Image	Semi-dense Image	Process Window	Material	EL (%)
<i>mask:</i> NTAR7 s_x^d, s_y^d : 60 nm, 88 nm s_x^s, s_y^s : 59 nm, 120 nm <i>source:</i> quadrupole σ_i, σ_o : 0.4, 0.8 α : 45 deg					15.8
<i>mask:</i> data base D : 109 nm s_x^d, s_y^d : 78 nm, 107 nm s_x^s, s_y^s : 70 nm, 134 nm <i>source:</i> quadrupole σ_i, σ_o : 0.4, 0.63 α : 29 deg					17.2
<i>mask:</i> free D : 67 nm s_x^d, s_y^d : 67 nm, 83 nm s_x^s, s_y^s : 52 nm, 114 nm <i>source:</i> quadrupole σ_i, σ_o : 0.19, 0.62 α : 30 deg					19.8

All cases yield solutions showing an adequate EL. Similarly as before, the flexible n/k range definition produces the best results, albeit employing non-available materials especially for the top absorber layer. In that case, also the absorber thickness is significantly lower in comparison with the restricted material selection mechanisms. Similarly to the projector-centric compensation of mask-induced phase effects (see Section 7.4), source configurations and mask biases appear to be closely linked to each other, thus suggesting the rigorous co-optimization of all components. This is substantiated by the fact that the respective source and mask settings in all of the cases differ considerably.

8.3 Mask Absorber Optimization for EUV lithography

The impact of the mask stack on the EUVL⁴ imaging performance has been a subject of research for more than decade. Bollepalli et al. (1998), for example, demonstrated a strong bias between features on dark and bright field masks. Moreover, because of the non-orthogonal angle of incidence on the mask, telecentricity errors are inherent in EUV lithography and lead to a significant orientation dependence of the pattern placement and the line widths. Since these errors are focus dependent, they induce focus shifts and thus lead to process window asymmetries (Krautschik et al., 2001).⁵ Other systematic studies by Schiavone et al. (2001) or Erdmann et al. (2003a) confirm these results, especially demonstrating the feature-dependence of these effects.

The aim of this study is to identify combinations of mask feature sizes, mask absorber parameters, and illumination settings to provide the best lithographic performance in terms of overlapping process windows and to reduce placement errors.⁶ Different feature types and AttPSMs and AltPSMs are investigated.

8.3.1 Optimization Variables

As with the previous source/mask optimization approaches, the optimization variables can be subdivided into two categories: (1) illumination parameters and (2) reticle-related parameters. The following paragraphs illustrate the specific representations used in this work.

Projection System

In this study, a 4x projection system with an NA of 0.3 and polarized light at a wavelength (λ) of 13.5 nm is employed. As usual, the angle of incidence is 6 degrees. Off-axis illumination has been selected, using an annular illuminator. The source geometry is varied by its geometric parameters, the inner and outer radius (σ_i/σ_o) of the annulus. In order to prevent the optimization routine from generating invalid solutions, a simple decoder strategy is employed: Instead of the outer and inner radii, the center radius (σ_c) and the latitude of the annulus (ℓ) are varied (as illustrated in Figure 8.9). The center radius is restricted to a range between 0.15 and 0.75, the annulus width ranges between 0.2 and 0.5. A granularity (delta) of 0.01 is used.

Absorber Layer Stack

All simulations are performed with a fixed multilayer stack consisting of 40 MoSi bi-layers. The molybdenum layer has a thickness of 2.74 nm and is assumed to have a refractive index (n) of 0.9211 and an extinction (k) of 0.0064 (both at a wavelength of 13.5 nm). The silicon layer is 4.11 nm thick; its optical properties are: $n = 0.99993$ and $k = 0.0018$ (at 13.5 nm λ). A silica (SiO_2) substrate is used. For simplicity, only a 11 nm bottom capping layer is introduced. No further capping or buffer layers are currently taken into account. However, their integration into the proposed approach can be done in a straightforward manner and is subject to future work.

⁴see Section 2.3

⁵Like in the optical case, the mask-induced phase variations can be interpreted as mask aberrations and hence formalized and quantified similarly as for the projector, using Zernike polynomials (Nakajima et al., 2008; Erdmann et al., 2010).

⁶The following results have been presented in (Fühner et al., 2008).




Figure 8.9: Source geometry optimization variables.

As the first optimization variable, the absorber's material is optimized, employing the same three alternative material modes as before:

- (1) *List of user-supplied stacks.* Here, only a single fixed standard chromium stack is provided. Its thickness was set to 80 nm, and the optical properties (at 13.5 nm) are assumed to be n, k : 0.9325, 0.0389.
- (2) *Material data base.* The materials used here are illustrated in Figure 8.10. Properties of the material indicated with white-filled circles (and black labels) are taken from Nugrowati (2008). The other materials (circles filled with black) are from Erdmann et al. (2003b).




Figure 8.10: List of materials and their corresponding refractive index (n) and extinction (k) at a wavelength of 13.5 nm.

- (3) *Free material properties.* The following restrictions were employed:

$$0.9 \leq n \leq 1.1$$

$$0 \leq k \leq 0.05$$

The optimization variables for the layer stack are demonstrated in Figure 8.11. They slightly differ between AttPSM and AltPSM. In both cases, the absorber thickness is varied within a range of 30 and 140 nm. The step width is 1 nm. Additionally, the line width is varied. The variation range strongly depends on the optimization goal, as shown in the results section (Section 8.3.3). Also depending on the optimization target, the pitch of the feature can be varied; alternatively, a bias between vertical and horizontal features can be optimized for.

In the AltPSM case, the phase shifter width and the number of etched bi-layers are additionally subject to variation. Furthermore, a bias between etched and un-etched openings can be used as optimization parameter.

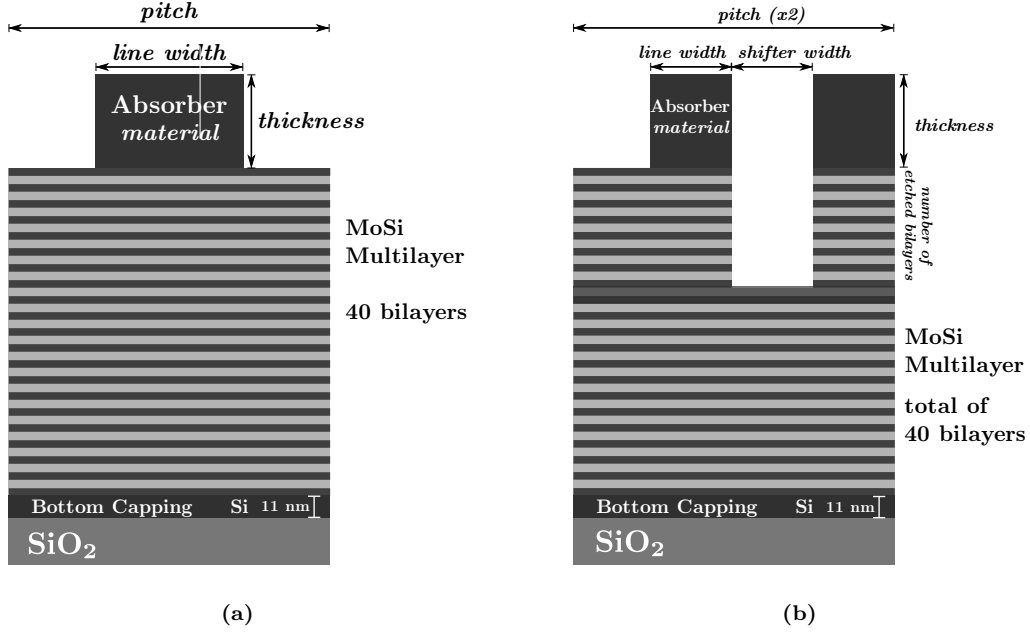


Figure 8.11: Mask layer stack. (a) AttPSM: variation of absorber thickness and material, line width, and pitch; (b) AltPSM: variation of absorber thickness and material, line width, shifter width, pitch, and number of etched bi-layers.

8.3.2 Optimization Criteria

Simulations are conducted with a full vector imaging approach without the Hopkins assumption. The mask is computed with the waveguide method, using a decomposition technique (Evanschitzky et al., 2007; Shao et al., 2008). A constant threshold model is used. Like in the previous studies, the main goal is the maximization of common process windows of different features, therefore employing the exposure latitude criteria shown in Section 8.1.2. In EUV lithography, because of the non-orthogonal angle of incidence, feature orientation and position dependent placement errors are a critical issue. Another important optimization objective is thus the minimization of such placement errors. In contrast to the approach proposed by Kang et al. (2008), it is assumed that a variation of the focus position is not a viable measure to counteract placement errors as it drastically degrades the contrast of images. Therefore the placement error is evaluated at the best focus position of the overlapping process windows. In addition to the two main objectives, process window overlap and placement error, several sub-criteria are introduced in order to stabilize and speed-up the convergence behavior of the optimization procedure. Depending on the specific optimization target, we have included evaluations of the image contrast and the NILS, the distance between the centers of individual process windows, and the absorber reflectivity and phase.

8.3.3 Results

Various results for different optimization goals are demonstrated. The image performance for line/space patterns is optimized using both AttPSM (incl. binary masks) and AltPSM. Additionally, in a first example, ideal AttPSM process conditions for contact holes are sought.

AttPSM: Line/Space Patterns

In order to evaluate the performance of a line/space patterning masks, two different goals are formulated: a through-pitch variation and a linearity evaluation.

Variation of pitch The following examples aim at the imaging performance for horizontal lines with a target width of 22 nm. For that, a through-pitch variation with duty ratios of 1:1, 1:1.2, 1:1.4, 1:1.6, 1:2, and 1:5 is evaluated. The main optimization criterion is to find a common process window for all duty ratios with the largest exposure latitude. The depth of focus (DOF) is fixed to 100 nm. Additionally, the placement error is to be minimized.

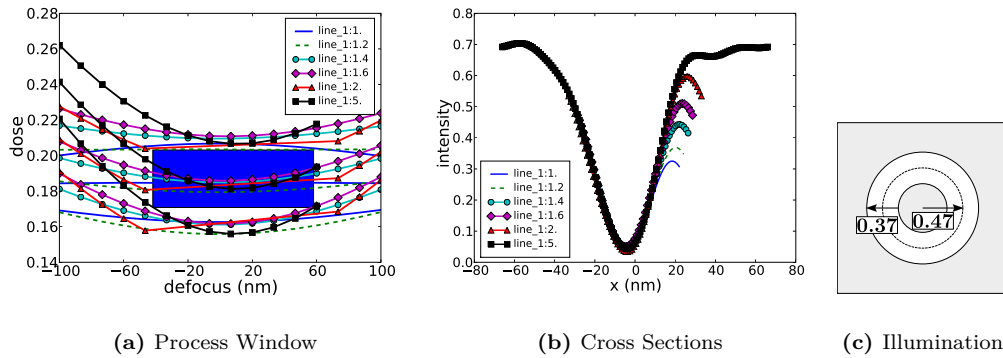


Figure 8.12: Process window, cross sections, and illumination settings for through-pitch optimization (AttPSM) solution with a fixed 80-nm chromium absorber.

The fixed 80 nm chromium absorber exhibits a reflectivity of 0.8 % and a phase of 117 degrees. With that, an exposure latitude of 17.3 % and a placement error of 4.34 nm can be achieved. The optimum line width on the mask for the different duty ratios are: 16, 17, 17, 17, and 19 nm; the annulus center has a radius of 0.47 and a latitude of 0.37. Figure 8.12 depicts the resulting process window and cross sections.




Figure 8.13: Process window, cross sections, and illumination settings for through-pitch optimization (AttPSM) solution using n/k list material representation.

The exposure latitude overlap can be drastically increased by optimizing both the absorber's thickness and its material. For example, with the data base representation, a 45 nm TaN absorber yields a weak

AttPSM (4.25 % reflectivity, -181 deg. phase) with a common exposure latitude of 22 % and a placement error of 3.42 nm (see Figure 8.13). The line configuration for this result is: 18, 19, 20, 20, 21, 17 nm.




Figure 8.14: Alternative solution for through-pitch optimization (AttPSM) solution using n/k list material representation.

An identical exposure latitude can be achieved with a high-transmission AttPSM (53.5 % reflectivity, phase: 154 deg.). In that case, a 48 nm thick absorber with a low extinction material (molybdenum) is found. The placement error is 3.84 nm. Line widths are: 16, 17, 17, 18, 19, 15 nm. Figure 8.14 illustrates this solution. Such a configuration can be expected to be very sensitive towards thickness variations, as will be shown in Section 8.3.3.

The more flexible approach with a free variation of the absorber's optical material properties generates solutions with slightly improved overlapping exposure latitudes and placement errors. For example, a low refraction, low extinction 39 nm absorber (reflectivity: 40 %, phase: -213 deg.) results in an exposure latitude of 26.7 % and placement error of 3.38 nm (see Figure 8.15). A similar performance is obtained with a high refraction and low extinction absorber (same thickness, reflectivity: 33 %, phase: -141 deg.). The exposure latitude in that case is 25 %, the placement error is 2.6 nm. As shown in Figure 8.16, in contrast to the previous low n solution, this candidate shows a process window shift toward negative defocus values. All solution candidates exhibit a significant process window asymmetry, primarily caused by the feature configuration with the largest pitch. Only moderate proximity effects are to be noted: The through-pitch line width varies by only 3–5 nm.

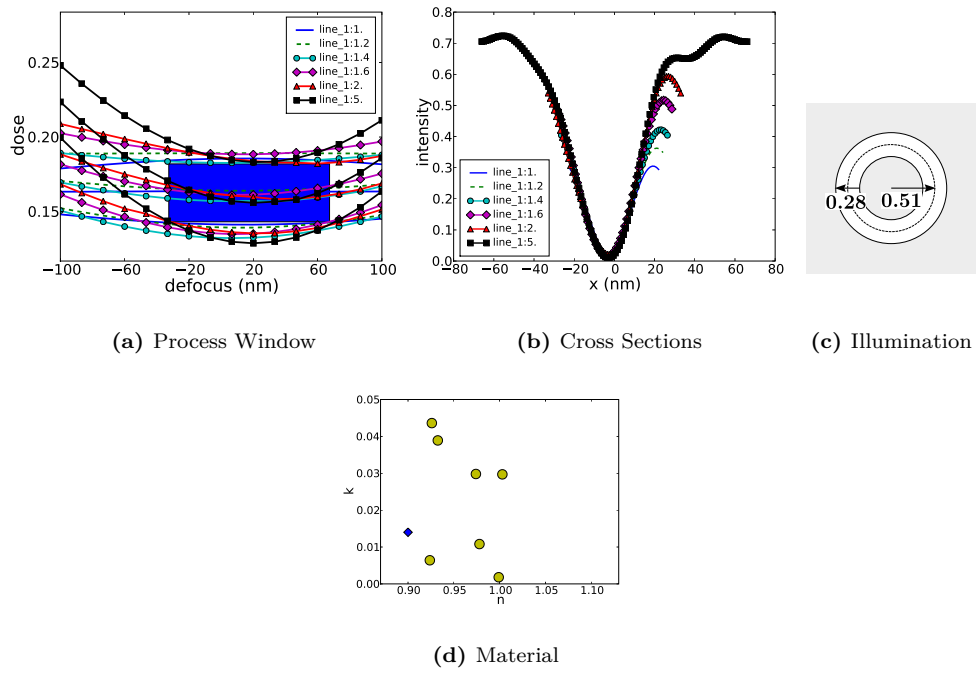


Figure 8.15: Solution candidate for through-pitch optimization (AttPSM) using n/k range representation. Refractive index is 0.9, extinction is 0.014.

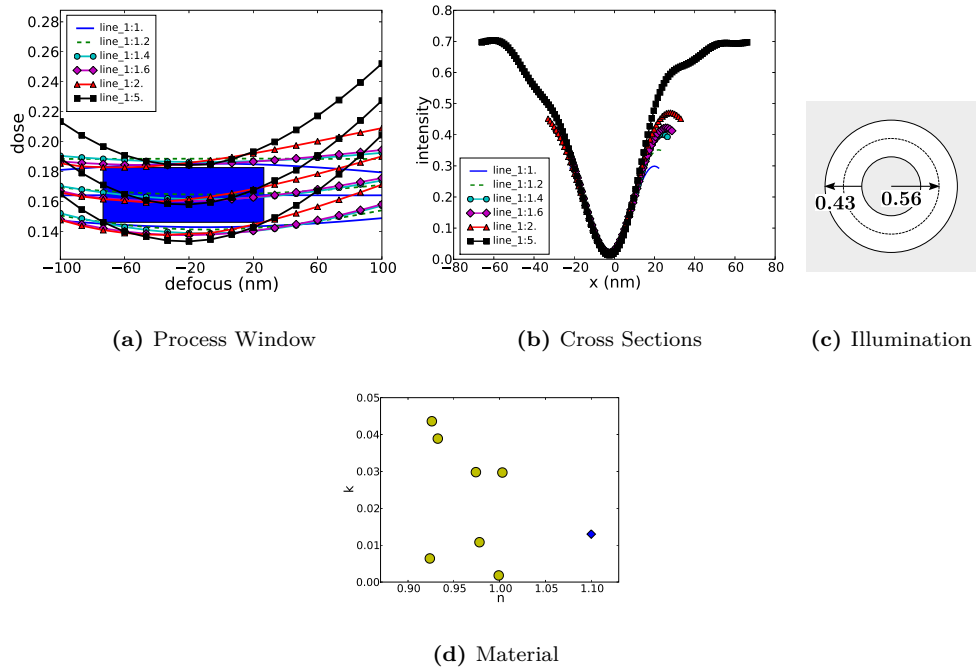


Figure 8.16: Alternative solution candidate for through-pitch optimization (AttPSM) using n/k range representation. Refractive index is 1.1, extinction is 0.013.

Table 8.4: Results for lines/spaces linearity optimization of AttPSM. In the first row, the solution's performance (reflectivity, phase, exposure latitude, and placement error) and the absorber thickness are listed; the second row shows the individual process windows for each target and the resulting overlap; cross sections are illustrated in row three; the last row shows the material properties of the absorber stack: Solution 1 uses a fixed chromium absorber. Solution 2 is optimized using the material list representation. Its absorber material (tantalum nitride) is highlighted in the material list plot. Solutions 3 and 4 employ the flexible n/k range representation; the obtained materials are indicated by the black diamonds. For comparison, the lighter circles depict the materials in the material list.

	Process Window	Cross Sections	Material
(1) refl.: 0.8 % phase: 117 deg. dose lat.: 16.35 % pl. error: 3.9 nm abs. th.: 80 nm			chromium
(2) refl.: 2.77 % phase: -203.2 deg. dose lat.: 18.73 % pl. error: 3.1 nm abs. th.: 52.0 nm			
(3) refl.: 8.79 % phase: -219.6 deg. dose lat.: 18.95 % pl. error: 3.5 nm abs. th.: 59.0 nm			
(4) refl.: 7.28 % phase: -133.43 deg. dose lat.: 18 % pl. error: 2.8 nm abs. th.: 42.0 nm			

Linearity In addition to a stable through-pitch behavior, an optimum source/mask combination should also show a near-linear characteristic for different target sizes. Therefore in the following examples, the optimization goal is to maximize the exposure latitude of overlapping process windows for the following line widths: 15, 22, 32, 45, and 65 nm. Again, the depth of focus is set to 100 nm.

The optimization results are presented in Table 8.4. The fixed 80 nm chromium absorber (solution 1) has a common exposure latitude of about 16%. The placement error is nearly 4 nm. With the material list representation (solution 2), solutions with an exposure latitude of almost 19% are found. The placement error is decreased to around 3 nm. The list material representation shows a notable preference of weak AttPSMs with high extinction materials such as chromium and tantalum nitride.

The absorbers' reflection values slightly increase when the material range representation is selected. Solution 3 uses a low- n /high- k material, whereas the material of solution 4 is characterized by both a high refractive index and a high extinction. In both examples, the extinction takes values close to the specified lower (0.9) or upper (1.1) bound, respectively. That causes the process window to be asymmetric either towards positive or negative defocus positions. In all solutions, especially the large features result in a distinctly shifted process window.

AltPSM: Horizontal and Vertical Lines

The following examples demonstrate first results of the proposed optimization procedure for alternating phase-shifting masks. The optimization goal is again to find ideal illumination and mask settings for line/space patterns. The number of etched bi-layers is varied in order to find an ideal phase difference. Moreover, to compensate for the intensity imbalance between etched and unetched openings, the shifter width is subject to variation. In this first example, the absorber's material is fixed to chromium. A fixed DOF of 300 nm is used. The target feature in this case is the imaging of 22 nm lines both in vertical and horizontal direction.

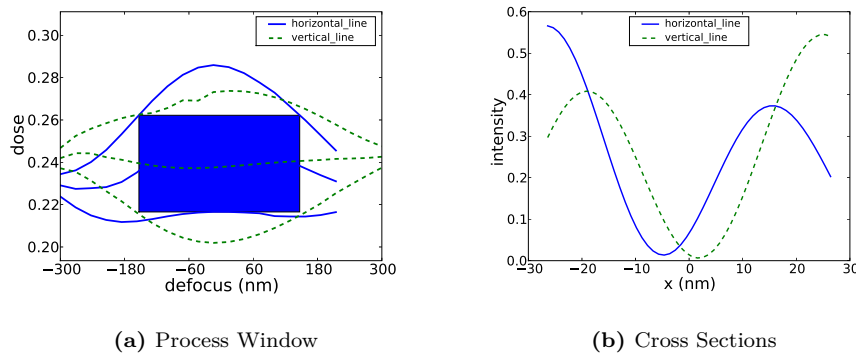


Figure 8.17: Solution candidate for AltPSM linearity optimization.

Figure 8.17 shows an example with a relatively thick absorber (117 nm). The horizontal line is 12 nm, the vertical line is 18 nm wide. Shifter widths are: 33 nm (horizontal) and 32 nm (vertical). With this configuration, an exposure latitude of 19% for the overlapping process windows and a placement error of 2.6 nm are achieved.

Another solution candidate with a similar performance (exposure latitude: 18.3%, placement error: 2.2 nm) but with a thickness of 76 nm is depicted in Figure 8.18. In that case, the horizontal line width is as small as 10 nm, the vertical line has a width of 15 nm. The shifter widths are almost identical to the previous example (horz./vert.: 33/31 nm). Both examples feature a coherent illumination configuration (center radius: 0.15, annulus latitude: 0.2), which is known to provide the best AltPSM imaging performance. Also, the number of etched layers, 13, is common to both cases.

Arrays of Contact Holes (AttPSM)

In another survey, 22 nm by 22 nm contact hole arrays are optimized. Again, the depth of focus is set to 100 nm. In order to evaluate the through-pitch performance, different duty ratios are taken into

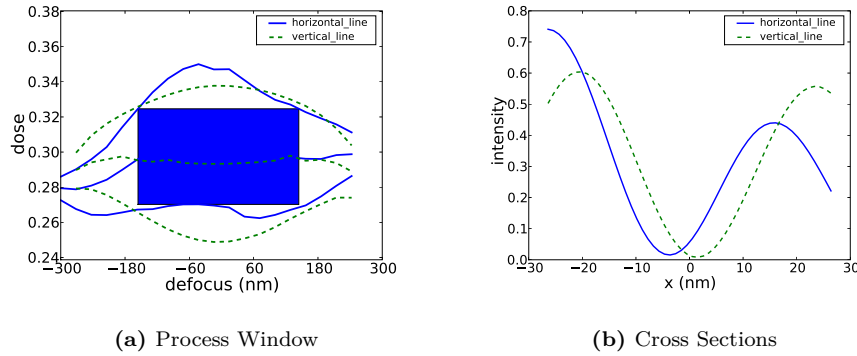


Figure 8.18: Alternative solution for AltPSM linearity optimization.

account: 1:1, 1:2, 1:5. Additionally, the dense (1:1) configuration is computed in horizontal and vertical direction (by switching between horizontal and vertical illumination). The CD difference between the two directions is counteracted by a supplementary optimization of the mask bias.

Table 8.5 summarizes the obtained results. The first example (1) serves as a reference solution. It uses a fixed 80 nm chromium absorber. Solutions 2 and 3 are optimized using the n/k material list representation. Both a low- and a high-transmission absorber are found. Another low-transmission solution with a thicker (114 nm) tantalum nitride absorber is also obtained: Its exposure latitude is nearly equal to that of solution 3; however, the placement error is slightly larger. Both solutions show low threshold intensities. Therefore they would require a high exposure dose. The high-transmission solutions, on the other hand, exhibit a contrast reversal, thus would have to make use of a negative photoresist.

With the more flexible n/k material range representation, similar solutions in terms of performance are found (see Table 8.5, (4) and (5)). Solution 4 depicts a low transmission mask with a very small placement error and a contrast reversal. The material in that case has both a low refractive index and a low extinction. Yet another well-performing solution (5) comprises a binary mask and with a low- n /high- k absorber.

Table 8.5: AttPSM contact hole optimization results. In the first row, process windows for the different duty ratios and the resulting overlap are shown; the second row shows the corresponding cross sections; row three depicts the illumination setting with the center radius (σ_c) and the annulus latitude (ℓ); the material property graphs in the last row correspond to the preset representation option (cf. Table 8.4): solution 1: fixed chromium absorber, solution 2: material list, solution 3 and 4: n/k range. Figures of merit and optimization variables results are listed underneath the graphs.

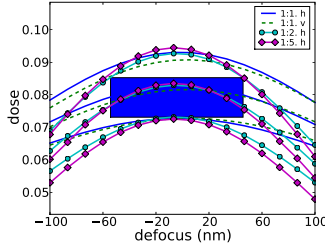
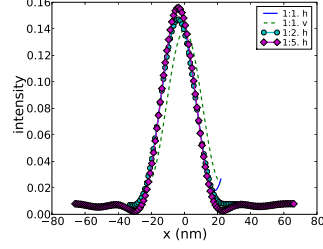
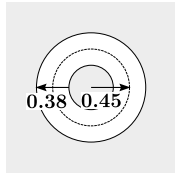
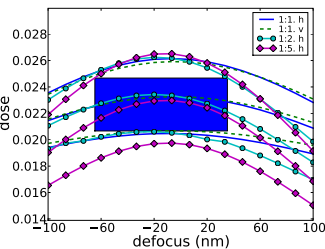
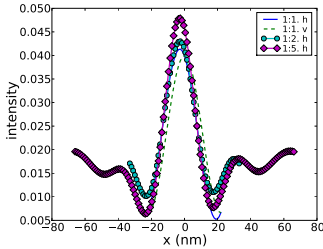
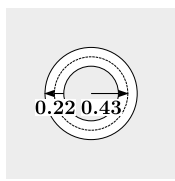
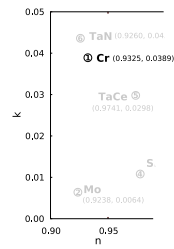
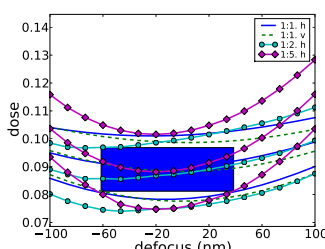
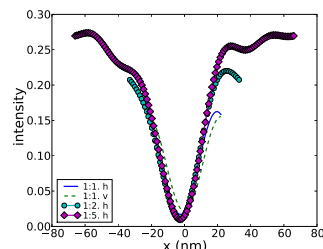
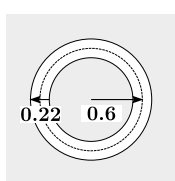
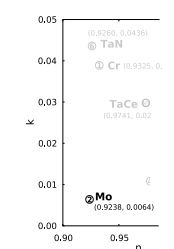
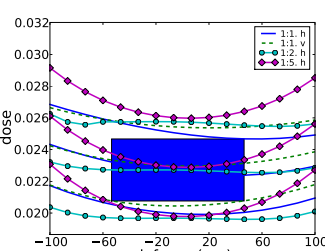
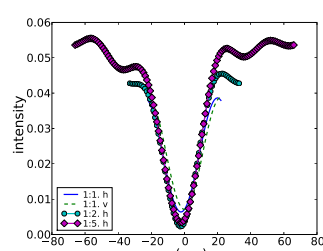
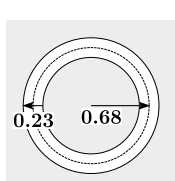
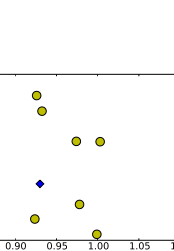
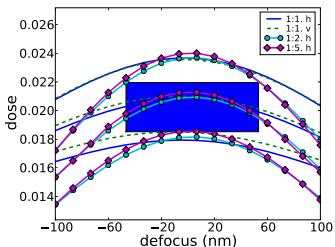
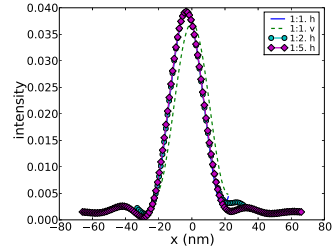
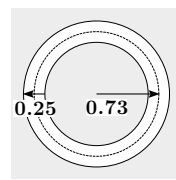
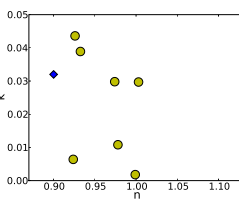
Process Window	Cross Sections	Illumination	Material
			chromium
(1) reflectivity: 0.8 %, phase: 117 deg., exposure latitude: 15.35 %, placement error: 2.8 nm line widths: 24, 24, 24 nm, horiz./vert. bias: 3, absorber thickness: 80 nm			
			
(2) reflectivity: 2.3 %, phase: 129 deg., exposure latitude: 17.67 %, placement error: 2.1 nm line widths: 16, 16, 16 nm, horiz./vert. bias: 8, absorber thickness: 60 nm			
			
(3) reflectivity: 40.91 %, phase: 92.6 deg., exposure latitude: 22.21 %, placement error: 2.13 nm line widths: 13, 15, 13 nm, horiz./vert. bias: 2, absorber thickness: 64 nm			
			
(4) reflectivity: 8.45 %, phase: -266.62 deg., exposure latitude: 22.32 %, placement error: 1.7 nm line widths: 13, 14, 13 nm, horiz./vert. bias: -4, absorber thickness: 72 nm			

Table 8.5: continued

Process Window	Cross Sections	Illumination	Material
			
(5) reflectivity: 0.17 %, phase: -53.18 deg., exposure latitude: 16.95 %, placement error: 2.45 nm line widths: 17, 17, 18 nm, horiz./vert. bias: 3, absorber thickness: 73 nm			

High Reflectivity Masks

As demonstrated with the previous results, highly reflecting absorbers show a good performance both in terms of overlapping process windows for different feature set-ups and placement error. Most absorber configurations for such high-reflectivity masks comprise materials with small refractive indices and extinction values, and a small thickness. However, with the proposed optimization approach only very few such absorbers were found. One could even observe that solution candidates with materials such as molybdenum became often “extinct” from the population at a very early stage of the optimization run. This hints at the fact that such configurations exhibit a high sensitivity against other optimization variables.

Figure 8.19 shows the dependence of a high reflectivity (Mo) absorber on its thickness. The optimization goal was to find an AttPSM with maximum overlapping process windows for dense and semi-dense lines in horizontal and vertical directions. The best solution with an exposure latitude of above 14 % exhibits an absorber thickness of 40 nm. Neighboring solutions with thickness values between 30 and 50 nm are evaluated. Within the range of ± 3 nm, the exposure latitude is still acceptable. The placement error hardly exceeds 2 nm. Process windows at different thickness positions (35, 37, 40, 43, and 45 nm) are portrayed. The left-hatched rectangles illustrate the process window overlap of the examined solution. Whereas the right-hatched rectangles indicate the overlap of the best solution. Already the 37 nm thick absorber has no process window in common with the 40 nm one. In other words, a thickness tolerance of as few as 3 nm would render this solution unfeasible.

8.4 Bottom-antireflective Coating Optimization

The previous considerations were solely concerned with the lithographic projection system, including mask and illuminator, for which we have shown that a rigorous optimization is suited to drastically increase its performance. In order for the improvements to be effective for the entire lithographic system, of course also the image target, the wafer, and its properties need to exhibit an adequate behavior. In this example,⁷ we aim at optimizing the optical properties of the wafer stack. There, one of the critical aspects is the reduction of back-reflection from the substrate, which due to the resulting standing waves is one source of pattern fidelity degradation. The reflectance from the substrate, which is in dependence of the material properties of the interfacing layers the angle and the polarization state of the incidence light, can be drastically reduced by introducing auxiliary layers. A common approach is the coating of the substrate with an anti-reflective thin-film, called BARC. In addition to exhibiting a high suppression of reflection under the targeted illumination conditions, the BARC layer and its deposition should have

⁷Summarizing the results from Liu et al. (2010).

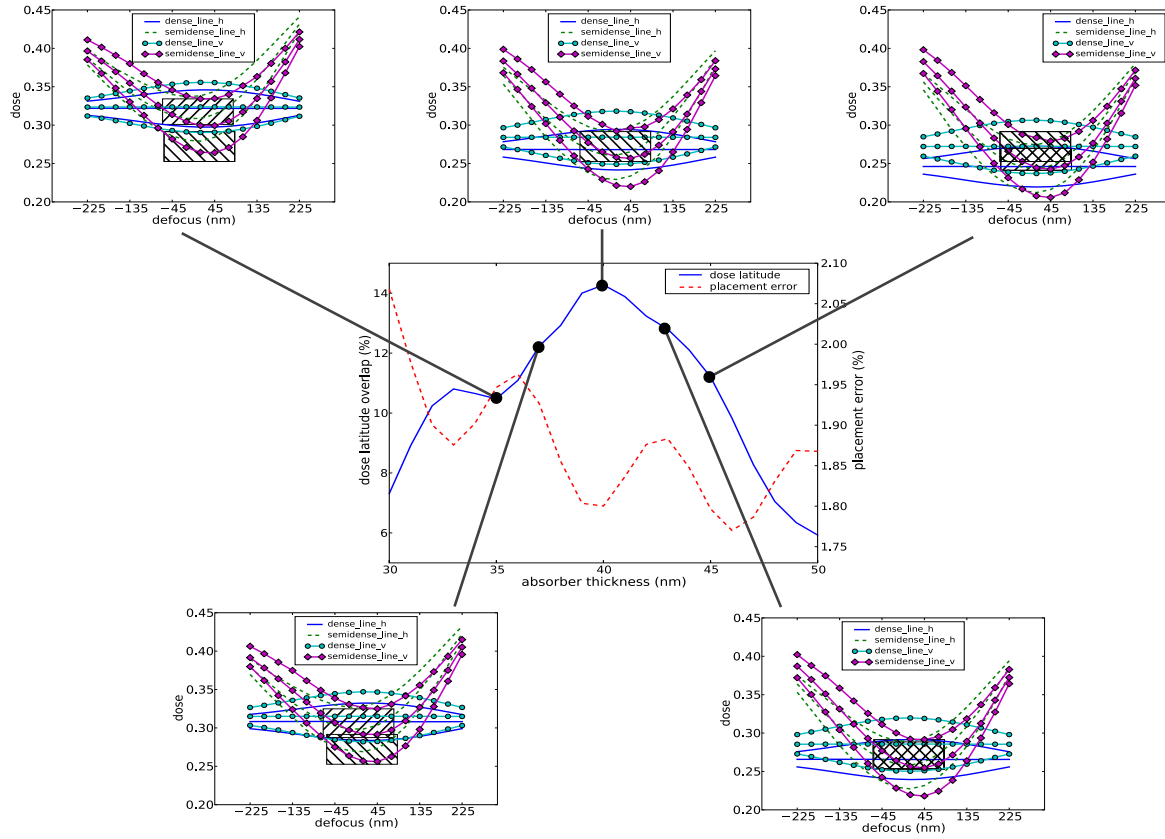


Figure 8.19: Sensitivity of highly reflecting absorber against thickness. The performance of the mask in terms of dose latitude and placement error strongly depends on the thickness of the highly reflective absorber. Only within a range of about ± 3 nm around the best, 40-nm thick solution an acceptable process window overlap is obtained.

as little impact on the manufacturing process steps as possible. This requirement is especially critical for double patterning techniques (see Section 2.6.4), which are prone to a severe throughput decrease due to the multiplication of steps involved. The use of organic BARC materials is therefore common practice, because they can be processed in the scanner track (Guerrero et al., 2008). In general, organic BARC materials are composed of polymer, cross-linker, cross-linking catalyst and solvent in order to meet different requirements such as chemical freezing, contamination blocking and reflection control at the same time. During the double patterning process, which this study is aimed at, the first lithographic exposure and subsequent process steps modify the properties of the BARC and impact its performance in the second lithography stage. For example, an exposure with a lines and spaces pattern can result in a laterally and sinusoidally modulated refractive index of the BARC. Experimental results showed that the difference of the index of refraction (n) before and after the first exposure and develop can be as large as -0.08 (Sakamoto et al., 2010). The change of the n value and its modulation in the BARC layer alters its reflectivity behavior in the second lithography step. Moreover, in immersion lithography systems, the increased NA leads to a significant polarization-dependent behavior of the wafer stack. Angles may in fact become as large as being close to Brewster's angle (Brunner et al., 2002). It therefore becomes increasingly difficult to identify BARCs with an adequate performance, even more in the case of double patterning processes.

8.4.1 Optimization Variables and Figure of Merit

The anti-reflective performance of a BARC is determined by the optical and topographical properties of its individual layers, modeled by the refractive index (n_j), the extinction coefficient (k_j) and the thickness (d_j) of each BARC layer $j = 1, 2$. These optimization variables and the corresponding wafer stack are depicted in Figure 8.20. Broadly, the optimization goal is given by the reflectance returned at the interface between the top BARC layer and the photoresist. In order to optimize the BARC performance for a large range of pitches of features on the photomask, the reflectivity (R) is evaluated for a number of different incidence angles (θ_i). The objective function for one lithography step (l) can hence, for example, be stated as the squared mean reflectance over a number of incidence angles (N):

$$\mu_l : \mathbf{x} \mapsto \frac{1}{N} \sum_{i=1}^N R_l^{\theta_i}(\mathbf{x}), \quad (8.2)$$

which is to be minimized. The vector \mathbf{x} denotes the optimization parameters:

$$\mathbf{x} = [n_1, k_1, n_2, k_2, d_1, d_2]^T.$$




Figure 8.20: Bi-layer BARC with optimization variables.

The assessment of the sensitivity of the BARC to the first lithographic exposure is performed under the simplified assumption of a sinusoidal modulation of the refractive index of the top layer of the BARC. A similar situation would arise from an exposure with a line/space pattern. The degree of index modulation is defined by a user-provided parameter Δn . The two-stages of the BARC assessment can be summarized as follows:

- (1) In the first lithography step, the layers can be assumed to be planar, allowing for the application of the time-efficient transfer-matrix method.
- (2) The second lithography step, however, exhibits an inhomogeneous BARC top layer so that a rigorous EMF computation has to be carried out, for which again the waveguide method is used.

Applying, for example, a weighted sum approach, the total objective can be stated as:

$$\mu : \mathbf{x} \mapsto \sum_{l=1,2} w_l \sqrt{\mu_l}, \quad (8.3)$$

where typically $w_1 + w_2 = 1$.

The computation time for 35 incidence angles (9, ..., 43 degrees) is less than 0.1 seconds for the planar case (litho 1) and less than 90 seconds for the second lithography step (using the waveguide method) on our computing cluster nodes (AMD Opteron 2352 @ 1 GHz and Intel Xeon E5405 @ 2 GHz).

8.4.2 Results

In this section, results obtained from different optimization studies are presented. First, optimization results for both a single and a bi-layer BARC are displayed and compared. Then, the impact of the refractive index modulation in the top BARC layer is shown, and several optimization studies for the second lithography step are presented. Finally, three different optimization studies co-optimizing the BARC for both lithography steps are demonstrated.

Optimization of single and bi-layer BARCs for the first lithography step

In a first approach, the BARC performance for the first lithography step with a single-layer BARC is optimized. The objective is to obtain a low reflectance within an incidence angle range from 9 to 43 degrees, which corresponds to a feature pitch of 360–83 nm. The figure of merit is given by the mean of the squared reflectivities over the angle range (8.2). The incidence medium is assumed to be a photoresist with a refractive index of 1.7 at 193 nm wavelength and the outgoing medium to be silicon with a complex refractive index of about $0.91 - 2.8j$. Since all layers can be assumed to be planar, the transfer-matrix method is employed (see Appendix D).⁸ Constraining the thickness of the BARC layer to 0–200 nm and its refractive index to 1–2.2, an optimum reflectance behavior is achieved at a BARC layer thickness of 23.4 nm and a refractive index of $2.05 - 0.57j$ (cf. Figure 8.21). The corresponding reflectivity curve (b) shows a strong pitch dependence: The minimum reflectivity is almost zero (at 29 degrees), and the maximum reflectivity is about 1.03 percent (at 43 degrees).

One of the main motives to apply BARCs is to achieve a high degree of pattern stability. For example, the critical dimension (CD) should exhibit only a low sensitivity to resist thickness variations. One way to assess the impact of the performance of the optimized BARC on the pattern stability is given by a so-called *swing curve* analysis,⁹ function of the thickness of the resist. Figure 8.22 shows this interrelation for the optimized single-layer BARC. The analysis is performed for the same pitches as used in the optimization. The target CD on the wafer is about 41.5 nm. Because of higher reflectivities at larger incident angles (i.e., smaller pitches), the wafer CD shows a non-linear relation to the resist thickness (Figure 8.22(a)). At the same time, due to the resist absorption, the wafer CD linearly changes with the resist thickness at smaller incidence angles (Figure 8.22(b)).

⁸The variant used here was developed by Liu et al. (2012).

⁹The term *swing curve* alludes to a typically sinusoidal nature of the CD change under resist thickness variations. This is due to the fact that standing waves are caused by interference of the incident and the reflected light in the resist, which is more or less constructive depending on the thickness of the resist.




Figure 8.21: Optimized single layer BARC: (a) reflectance performance, (b) optimum solution properties.

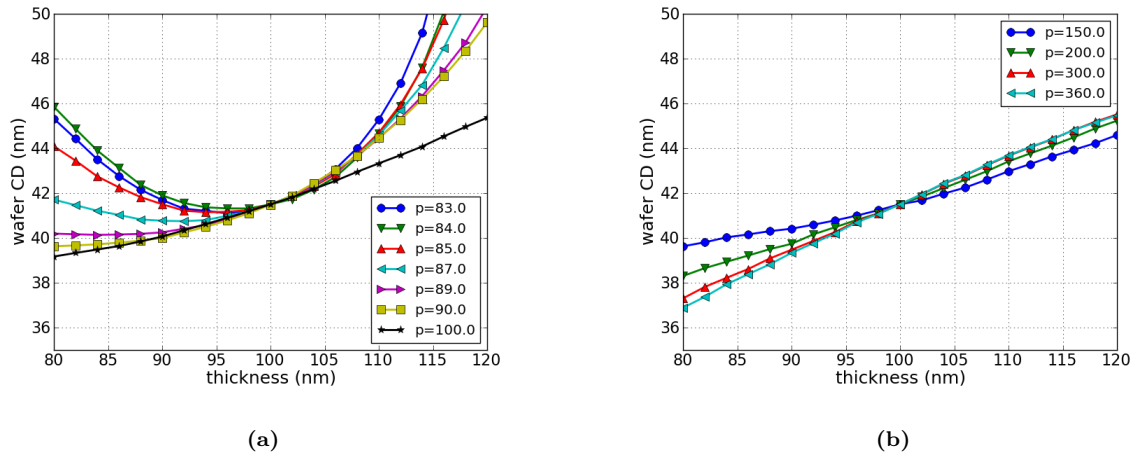


Figure 8.22: Pattern stability achieved with a single-layer BARC optimized for the first lithography step. The target CD is 41.5 nm: (a) For small pitches, a non-linear CD dependence on the resist thickness is observed, caused by the decreased antireflection performance of the BARC at higher incidence angles; (b) larger pitches show a resist absorption-induced linear behavior.

In a second optimization step, a bi-layer BARC is optimized. Again, the goal is to minimize the reflectivity within the incidence angle range from 9 to 43 degrees. The same thickness and optical property constraints as in the former optimization task are applied—this time, however, to each of the two layers. Moreover, the same merit function as before is used.

The genetic algorithm is run for 800 generations. However, after about 300 generations a local minimum that closely resembles the final solution is obtained. As shown in Figure 8.23(b), the optimum bi-layer stack consists of a top layer with a thickness of 13.2 nm and a refractive index of $2.08 - 0.13j$. The second layer has a thickness of 46.7 nm and a refractive index of $1.675 - 0.633j$. The reflectivity performance is drastically improved compared to the best single layer BARC (Figure 8.23(a)). The maximum reflectance, which occurs at an angle of 43 degrees, is less than 0.04 percent. In contrast to the single layer case, no significant pitch-dependence occurs, not even for small pitches.

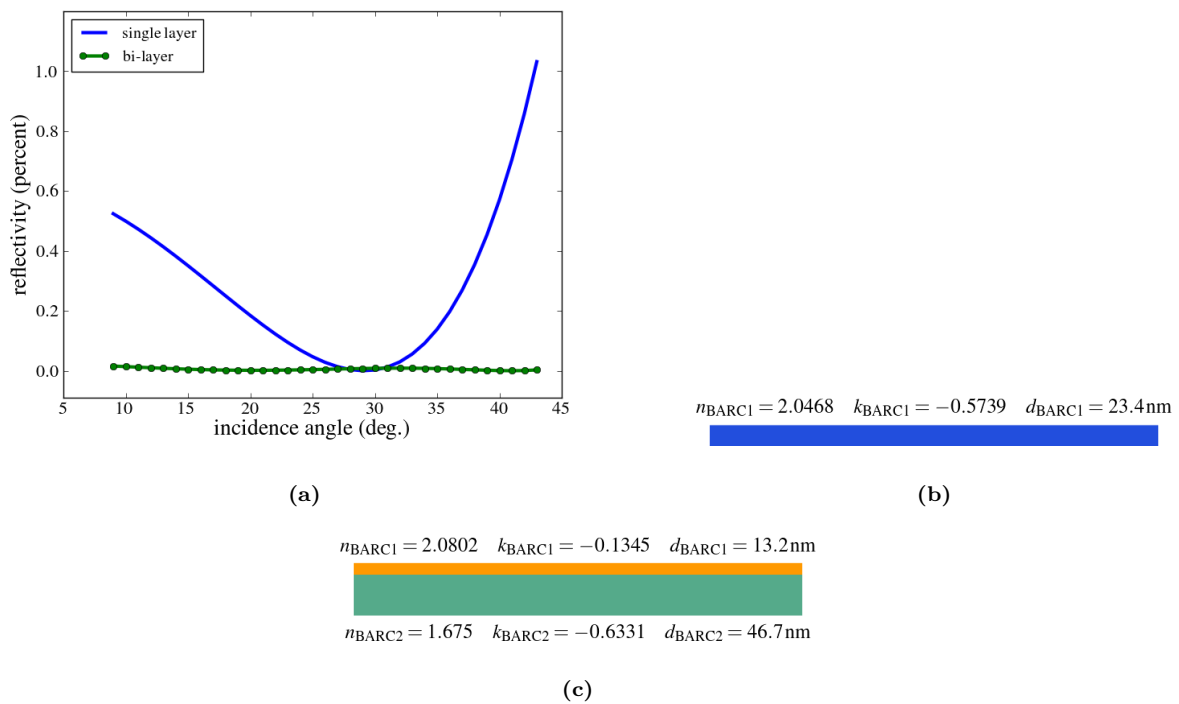


Figure 8.23: Comparison of the optimized single and bi-layer BARCs. (a) Anti-reflective performance within the incidence angle range of [9,43] degrees; thickness and optical properties of the optimized (b) single and (c) bi-layer BARC stacks.

The CD swing curves obtained with the optimized bi-layer BARC is demonstrated in Figure 8.24. The target CD on the wafer is about 45.1 nm. The swing curve shows an almost perfectly linear behavior in the incidence angle range of concern, allowing for an adequate control of the CD values on the wafer.

Bi-layer BARC with index modulation

After the first lithography step, the optical properties of the top layer of the BARC are modified due to both the exposure itself and subsequent process steps including baking. The topography of this index modulation can be assumed to be related to the patterns imaged in the first lithography step. For example, line/space (L/S) patterns will result in a sinusoidal modulation of the top layer index, whose period of this modulation is determined by the pitch of the L/S features. Figure 8.25 schematically demonstrates the resulted topography.




Figure 8.24: CD swing curves obtained with the optimized bi-layer BARC: for both (a) small and (b) large pitches an almost perfectly linear behavior is achieved.




Figure 8.25: Structure of bi-layer BARC with index modulation in one period

In order to assess the impact of the refractive index modulation, a simulation study is conducted in which the influence of an assumed index modulation factor in dependence of the pitch is investigated. In the modulated region, the refractive index is given by the sum of the original refractive index and the modulation factor (Δn). Both the BARC and the resist properties are taken from the former optimized bi-layer BARC (cf. Figure 8.23): A refractive index of $2.08 - 0.135j$ and a thickness of 13.2 nm are assumed for the top layer, and for the second layer, a refractive index of $1.68 - 0.633j$ and a thickness of 46.7 nm are applied. Index modulation values in the range of -0.5 to 0.5 are investigated, giving a refractive index variation of the first layer from 1.58 to 2.58. As before, the range of incidence angles is 9–43 degrees.

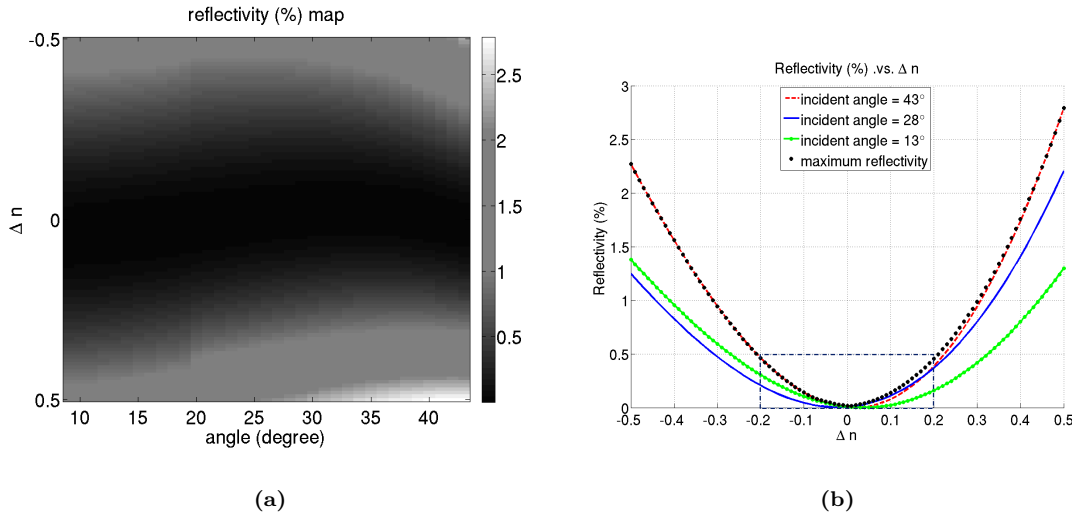


Figure 8.26: Reflectivity of BARC in lithography step two (a) as a function of index modulation of top layer and incidence angle. (b) Cross-section for selected incidence angles.

The resulting bi-layer BARC performance in the second lithography step as a function of the incidence angle and the index modulation value is shown in Figure 8.26(a). The period of the sinusoidally modulated region in the top layer is determined by the angle of incidence, i.e., the pitch of a given line/space feature (see Figure 8.25); only one period is evaluated. For small pitches (that is, large angles of incidence), the modulation impact is significantly larger than for large pitches. Figure 8.26(b) depicts corresponding cross-sections for selected angles.

In order to account for the effects of index modulation resulted from litho 1, the bi-layer BARC is next optimized for the second lithography step. Simulations for typical settings have shown that most of the energy of the reflected light is contained in the zeroth and minus first diffraction order. We have therefore modified the objective function (0) by replacing the reflectance by the average reflectance of both orders. In addition to the same constraints as applied in the previous optimization study, an index modulation (Δn) of 0.2 was assumed. As with the modulation impact study, the period of the modulated region is determined by the pitch (incidence angle) under consideration. The same incidence angle range of [9, 43] degrees is applied.

After 500 generations, an optimum solution was obtained—a very similar solution is already achievable after about 250 generations. The corresponding stack showing all optical properties and the thickness of the two BARC layers is illustrated in Figure 8.27(a). A significantly thinner top layer than in the BARC optimized for litho 1 (Figure 8.23) is obtained. The optical properties of the layer are almost identical to the litho 1 optimization step. The second BARC layer shows a considerably larger thickness but a smaller extinction coefficient. In Figure 8.27(b), the reflectivity performance in litho 2 for both the stack optimized for litho 1 and the stack optimized for litho 2 is compared. For both stacks, an index modulation value of 0.2 is assumed. An index modulation of 0.2 in the BARC top layer has significantly

increased the total reflectivity of the wafer stack, as illustrated by the dashed line in Figure 8.27. The maximum reflectivity is about 0.54 percent at 38 degrees. If the index modulation is considered in the optimization process, the total reflectivity can be greatly decreased, as illustrated by the solid line in Figure 8.27: The maximum reflectivity is reduced to about 0.066 percent.

Assuming a negative modulation value (in this example: -0.2) yields a slightly different optimum BARC (see Figure 8.27(c)). Comparing the litho 2 performance—under the assumption of a negative modulation value—of both the litho 1 and the litho 2 optimal BARC shows a qualitatively comparable difference as in the positive modulation value case (Figure 8.27(d)).

Co-optimization of BARCs for both lithography steps

As shown above, BARCs optimized for only the first lithographic process step cannot be guaranteed to show an adequate performance also in the second lithography step. Thus, in order to identify an appropriate bi-layer BARC for double patterning applications, a co-optimization taking both litho steps into account has to be performed. For that purpose, we have combined both evaluation routines—for litho 1 using TMM and for litho 2 using the waveguide method—in a weighted sum approach (see (8.3)). As in the litho 2 only case, the average of the two contributing orders (-1 and 0) is assessed. A refractive index modulation of 0.2 is assumed.

After 800 generations using the single-objective genetic algorithm, a bi-layer BARC with an optimum anti-reflection behavior in both lithography steps is obtained; a solution that closely resembles the final one is already achieved after 400 generations. The optimized stack and its performance are portrayed in Figure 8.28. The maximum reflectivity in litho 1 is about 0.1 percent at an incidence angle of 43 degrees and 0.13 percent at 9 degrees in litho 2. Again, the resulted top layer is extremely thin, jeopardizing the manufacturability of the BARC and suggesting the introduction of a feasibility criterion.

Bi-objective bottom layer optimization

In the previous simulation studies, high index modulations of 0.2 or -0.2 were assumed. Related experimental studies by Sakamoto et al. (2010) suggest smaller maximum modulation values. In a further optimization study, we have therefore applied BARC top layer properties reported by the authors: A refractive index of $1.8 - 0.3j$ and an index modulation factor of -0.08. In addition, the thickness of the top layer is fixed to 30 nm, also resolving the manufacturability issue resulting from the extremely thin top layers obtained in the previous study. Both the thickness and the optical properties of the second BARC layer are subject to optimization. The thickness range is set to [10, 200] nm, the refractive index was optimized in a range of [1, 2.2] nm, and the extinction coefficient between 0.1 and 0.5. A bi-objective optimization approach using a multi-objective genetic algorithm (SPEA2 with an unlimited archive) was applied for the optimization. All other settings, including the number of incidence angles was unaltered compared to the previous studies.

Figure 8.29 illustrates the solutions obtained with the bi-objective optimization approach. The Pareto front, depicted in the center of the figure, forms the optimality set. That is, solutions on the Pareto front outperform other candidates either with respect to litho 1 (μ_1) or litho 2 (μ_2), while still showing an acceptable performance in the respective other litho step. Solutions with a perfect compromise between litho 1 and litho 2 can be found close to where the Pareto front intersects the bisectrix of the coordinate system. The circumjacent graphs demonstrate the anti-reflective performance of example solutions taken from the Pareto front. Despite the reduced index modulation, the performance for both litho 1 and litho 2 is significantly degraded compared with the previous study in which also the top layer of the BARC was subject to optimization. Solutions optimal in litho 1 (e.g., solution 1237) have a bottom BARC layer with a extensively higher index of refraction than solutions optimal in litho 2 (e.g., solution 3214). The two example compromise solutions (1551 and 2207) corroborate this tendency: The closer solutions are situated to the minimum of μ_2 , the lower is the refractive index of the bottom layer. This behavior can be attributed to the fact that higher refractive indices in the second layer of the BARC may increase back-reflections from the bottom layer into the top layer and hence increase the sensitivity of the solution to the index modulation of the top layer, which plays a role only in litho 2.

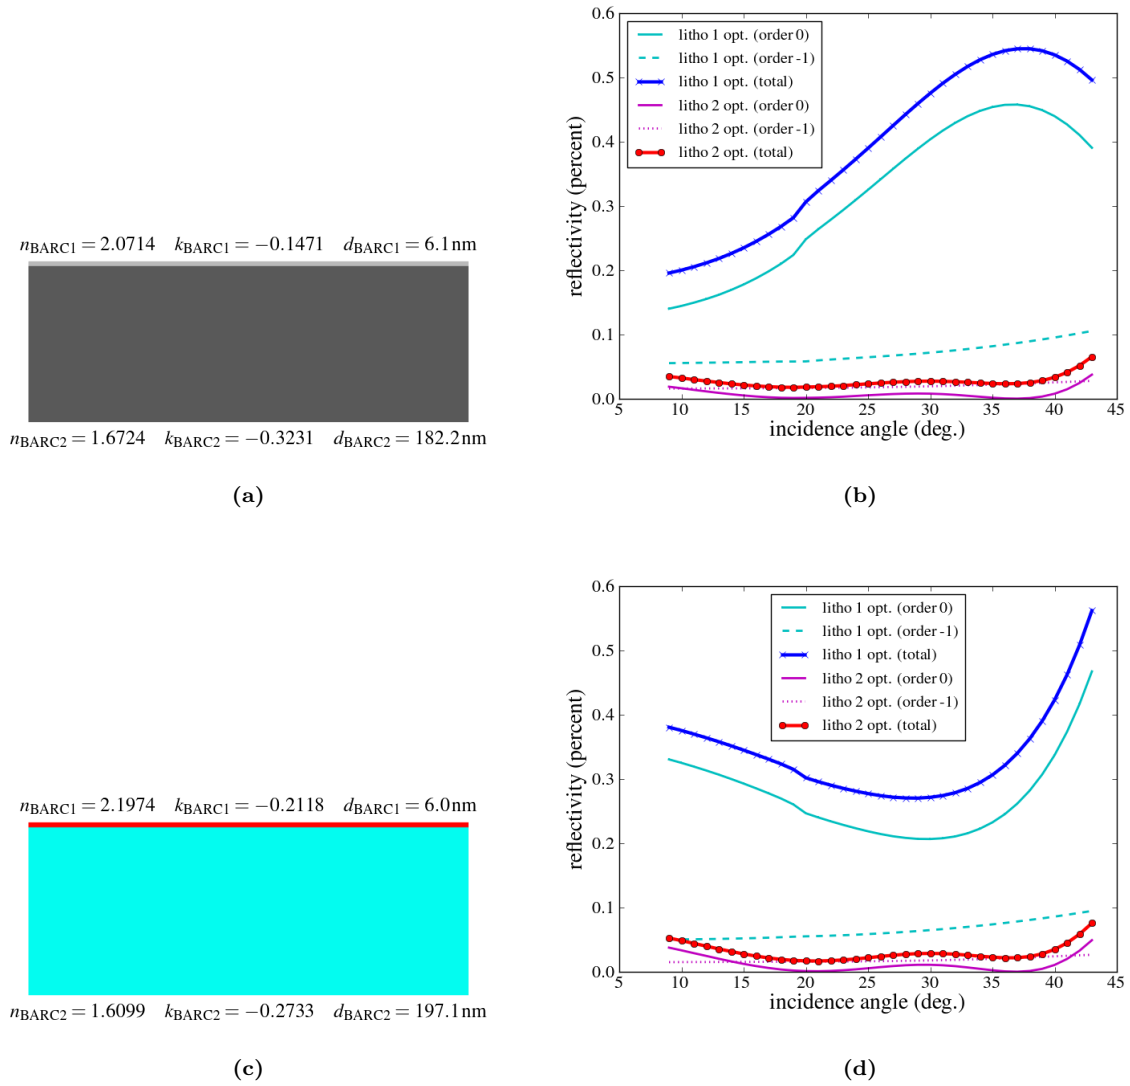


Figure 8.27: Bi-layer BARC optimized for the second lithography step, assuming a sinusoidal index modulation of 0.2. (a) Resulting stack; (b) comparison of litho 2 performance (total reflectivity of both contributing orders) of stack optimized for litho 2 (line marked with dots) and stack optimized for litho 1 (line marked with crosses). The other lines show the reflectance of the individual minus first and second orders; (c) stack optimized for negative index modulation, (d) evaluation of litho 1 and litho 2 optimized stacks in the negative modulation regime




Figure 8.28: Bi-layer BARC co-optimized for litho 1 and litho 2: (a) Anti-reflective performance, (b) geometry and optical properties of the stack.




Figure 8.29: Pareto front (center) and example solutions obtained with a bi-objective optimization study: The figures of merit are constituted by the two lithography steps (litho 1: μ_1 , litho 2: μ_2). For the second lithography step, the reflectivity of the minus first, the zeroth and the total reflectivity are displayed. Solutions performing well in the first lithography step (e.g., 1237) show a high refractive index of the bottom layer of the BARC. Solutions with a low reflectivity in step two (e.g., 3214) exhibit a higher index of refraction in the bottom layer. Compromise solutions can be found where the Pareto front is close to the bisectrix of the coordinate system, for example 1551 and 2207.

For the same reason, the extinction coefficient in all solutions is set to the maximum value of 0.5. No significant variation of the thickness of the second BARC layer in the Pareto set can be observed.

Tri-objective Bottom Layer Optimization With Variable Top Layer Thickness

In order to further improve the performance of the bi-layer BARC, we have carried out another optimization study, allowing for an additional degree of freedom by subjecting the thickness of the top layer to optimization. In contrast to the single-objective approach demonstrated in Section 8.4.2, the thickness of the first layer is constrained to $[10, 200]$ nm. To strictly account for the feasibility of solution candidates but also to investigate the impact of the top layer thickness on the performance of BARCs, we have also incorporated an additional thickness penalty:

$$\begin{aligned} \min \quad & p_d = d_1^u - d_1 \\ \text{s.t.} \quad & d_1 \leq d_1^u, \end{aligned} \tag{8.4}$$

where d_1^u is an upper thickness bound (here, 200 nm) of the first BARC layer. Hence by introducing a penalty like (8.4), a BARCs with a thicker top layer are preferred.

Figure 8.30 shows the Pareto set and example solutions obtained after 1000 iterations. As before, the X and the Y axes denote the figures of merit in litho 1 and litho 2, respectively. The thickness penalty criterion is constituted by the Z axis and also by the gray levels of the solution points. Depending on the preferred figure of merit, both the layer thicknesses and the optical properties of the second BARC layer of the optimum solution candidates show significantly different values. Solutions with a thick top BARC layer show a relatively thin second layer. All solutions optimal in litho 2 (μ_2) exhibit thinner BARC top layers.

As observed in the previous studies, thin top layers lead to a decrease of reflectance in the second lithography step due to a lower sensitivity to the refractive index modulation. In addition, an inverse effect can also be observed: Thin bottom layers lead to lower diffraction efficiencies in the modulated top layer, which in turn also decreases the impact of modulation and thus results in smaller reflectance differences between litho 1 and litho 2—however, leading to higher absolute reflectance values. Especially, BARCs with a top layer thickness close to the maximum of 200 nm (e.g., candidate 1537), where the impact of the bottom BARC layer is negligible, exhibit a poor reflectance performance at large angles. Solution candidates with a medium top layer thickness (e.g., 630, 727 and 1970) show an increased impact of the second BARC layer on the overall anti-reflective performance.

Candidate 3025 constitutes a feasible solution both in terms of manufacturability (27.4 nm thickness of top layer and 72.3 nm thickness of the bottom layer) and reflectivity performance (below one percent). In addition, the difference between litho 1 and litho 2 is small both absolutely and relatively, since both total reflectance curves show a similar behavior. Furthermore, the angle-dependent difference is small compared to all other results.

For top layers with a thickness of 20 nm and below (e.g., 63, 1503, 3773), the reflectivity of both lithography steps drops to below 0.5 percent. The difference between litho 1 and litho 2 is extremely small on the entire angle range in terms of absolute values. There is, however, a stronger reflectance dependence on the incidence angle than in the previous solution (3025). The manufacturability of these solutions may not be guaranteed.

One should note that the materials presented in this study are purely hypothetical. In order to further improve the feasibility of the proposed solutions, a look-up mechanism for existing materials, for example, similar to the one in the previous study, would have to be employed.

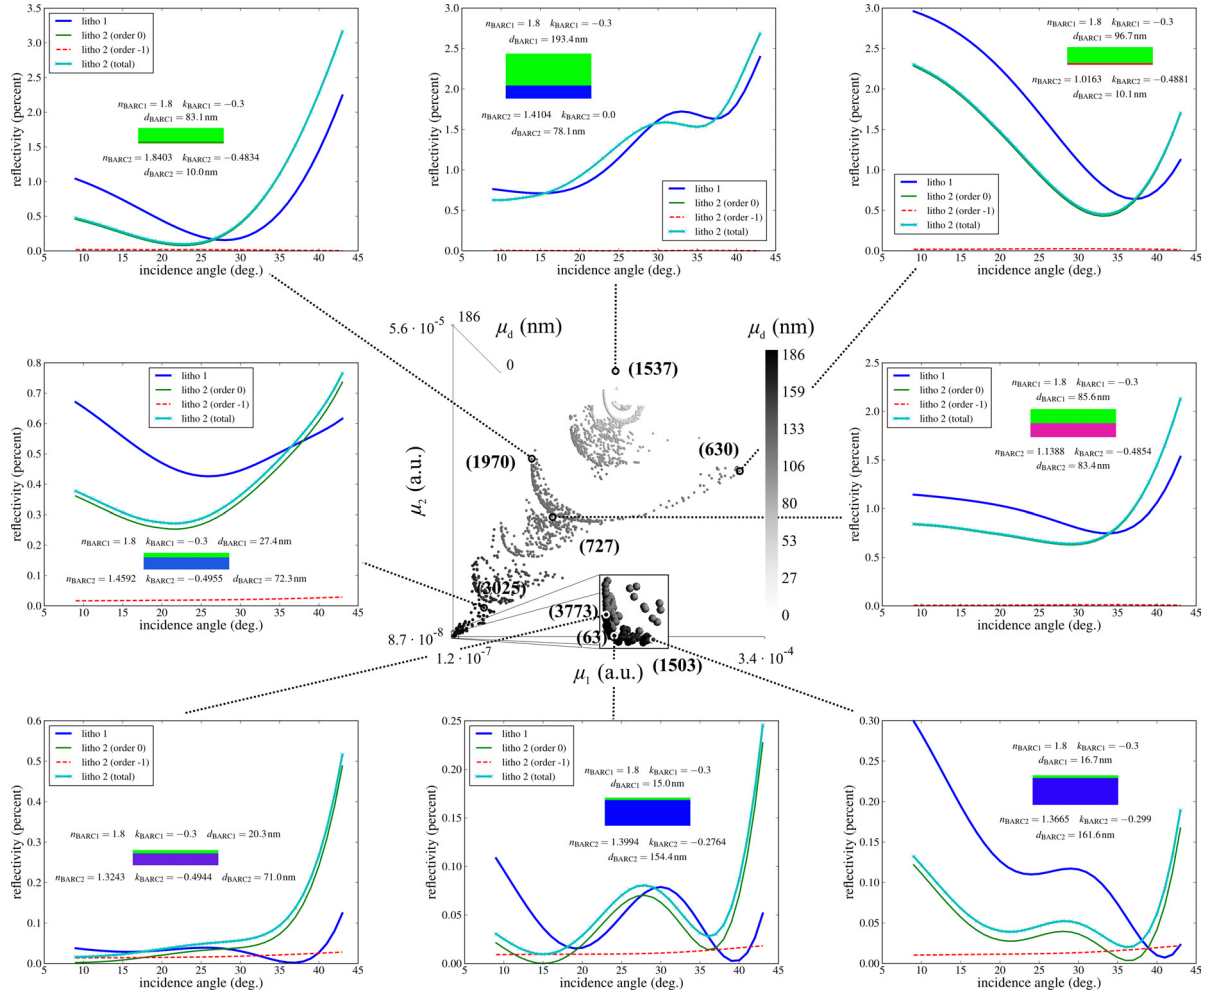


Figure 8.30: Different bi-layer BARC solutions obtained with the multi-objective optimization approach. In the center, the three dimensional Pareto front is shown. The X and Y axes denote the objective of litho 1 and litho 2, respectively. The Z direction and the gray levels of the solution points indicate the thickness penalty of the first BARC layer. The anti-reflection performance of the example solution candidates are displayed by their stacks, the total reflectivity of the first lithography step, and the reflectance of the zeroth and the minus first order, and the total reflectance of litho 2. The depicted solutions are taken from different regions of the Pareto front. For example, solution 1537 exhibits a large top BARC layer. However, its reflectivity performance is worse compared with solutions having thinner top layers (e.g., 3773, 63 and 1503). Compromise solutions can be found in the center of the Pareto set (e.g., 727).

9 Conclusions and Future Work

In this thesis, we have developed and demonstrated a novel, direct optimization approach for a wide range of problems related to lithographic process conditions. Through a large number of numerical experiments, we have shown in depth that this method allows to incorporate virtually any model options, evaluation mechanisms and constraint handling strategies by posing the optimization tasks as blackbox problems. The problems were solved with population-based algorithms that mimic evolutionary operations such as selection, recombination and mutation. They are hence commonly termed evolutionary algorithms (EAs). EAs can be considered global optimization routines in two respects, (1) they explore the entire search space and do not depend on local characteristics, (2) they globally recombine solutions that exhibit an above-average fitness. We have shown that these properties also prove beneficial in the context of this work. For example, in our approach to a rigorous, simultaneous optimization of the mask layout and the illumination source, a problem of increasing importance in computational lithography and electronic design automation (EDA), a variety of innovative solutions were obtained fully automatically. In contrast to related approaches, an actual co-optimization of the source and the mask was achieved. No nested optimization loops or post-optimization procedures were needed. In addition to a competent optimization algorithm, such a procedure depends on a stringent formulation of the problem, especially the parameter representation, evaluation routines and constraint handling mechanisms. We have shown that a direct, for instance, rectilinear representation of the mask is both feasible and suited to counteract a number of issues that occur under related approaches.

A number of authors have proposed to perform the mask optimization in the frequency domain. In that case, the optimization variables correspond to the amplitude of the diffraction orders. This approach can of course be considered elegant since it drastically reduces the number of parameters. The obtained result, the ideal assignment of the diffraction orders, however, does not have a unique spatial domain correspondence. To solve this ill-posed problem, another optimization step is required. The search for a Kirchhoff mask yielding a spectrum that closely resembles the target spectrum can be performed in a fairly straightforward manner. Accounting for electro-magnetic field (EMF) effects, especially considering incidence-angle-specific spectra, is far more involved. In both cases, the obtained mask layouts are in general not feasible, for example, since they are not manufacturable. Thus constraints or regularization routines have to be applied. Their result, in turn, may be quite different to situations in which the constraints are directly employed in the first optimization step. In fact, the decomposition into two steps can be generally expected to yield suboptimal results.

Others have used a spatial pixelated representation, which partly alleviates the former problem. Nonetheless, an enormous effort still has to be invested to prohibit a premature convergence toward non-feasible masks. This is accomplished by penalty or regularization methods, which are prone to exhibiting an unintentional bias. The manufacturability criterion in our approach, in contrast, does not constitute a means to resolve problems related to the non-convexity of the search space. It is solely integrated with a view to directly reflect the capabilities of the mask making process. We have demonstrated two alternative producibility criteria: A straightforward geometric one, which evaluates the layout of the structures, and a morphological one, which provides a measure for the integrity of structures under imprecisions as they may occur during mask writing. Since the evolutionary algorithm used in this work is specifically tailored to maintain a high level of diversity through what is commonly referred to as niching techniques, the optimization procedure faces little difficulties dealing with the multimodality of the search space. On the contrary, the set of candidates obtained through our optimization procedure typically contains various solutions that can be considered local optima. That is, in addition to the quasi-global optimum, the decision maker can optionally review the set of solutions, many of which present similarly appropriate performances as the champion.

The illumination source representation in this work follows a correspondingly straightforward strategy as the mask: Instead of the commonly applied pixelation on a Cartesian grid or an archel-based representation, which is well-suited for a simplified frequency domain analysis but which does not provide an adequate flexibility, for example, when considering incidence-angle-specific diffraction spectra, we propose the use of a polar-coordinate system. With this approach, the optimization accuracy can be chosen such that it directly corresponds to the accuracy of the imaging algorithm; so avoiding both a numerically non-significant overdetermination of source variables and an accuracy loss due to too few source points. We have shown that such an approach is also well suited to account for the technical prerequisites of the illumination source: The track/sector representation closely resembles the traditional use of aperture stops, while a source-point-based representation is eligible for state-of-the-art sources using diffractive optical elements or mirrors.

Most of the related approaches use a target image in order to minimize its deviation from images obtained for the candidate solutions, claiming to so solve an ill-posed inverse problem. This is problematic in a number of respects: Generating a suited target image is far from being straightforward. For example, a binary representation does not provide a suited goal since it is unachievable and hence provides very little control of the actual image performance. Using weights to accentuate critical or attenuate non-critical layout regions only solves this problem to some extent. Specifically, the a priori selection of appropriate weights is tedious and does not guarantee for solutions actually complying with the decision makers intentions. The composition of more suited images incorporating smooth slopes, on the other hand, implicates a “search direction” that may even result suboptimal solutions compared to potential solutions that meet the process requirements but are not restricted to an artificial target.

In our work, we have hence not employed any such target pattern comparison. We have instead relied on CD measurements and have shown that such an approach is sufficient to meet size requirements up to virtually any desired precision. But we have only implicitly considered an edge placement fidelity, whose more rigorous inclusion would be a requirement for future extensions toward an EDA solution. An according criterion can be facily incorporated but can also be expected to slightly increase the convergence time. A number of additional optimization criteria have been developed and their feasibility was shown. For example, in our approach, local contrast criteria such as normalized image log slope (NILS) can be directly incorporated. Moreover, when using sub-resolution assists, sidelobes have to be prevented from printing. In most of the other approaches, the incorporation of such a criterion is far from being apparent and can only be indirectly formulated, for instance, by minimizing the intensity of dark regions. Such an approach is problematic in that it depends on a strict difference metric (infinity norm) to capture local deviations from the intended intensity, which, however, poses a problem for the majority of gradient-based optimization approaches. Moreover, albeit a combination of constraining bright regions to a predefined minimum intensity and minimizing the intensity in dark areas is suited to increase dose latitude, it artificially constricts the set of potentially equally well-performing solutions. Similarly, all of the related approaches show a certain difficulty in accounting for the through-focus behavior. In our procedure, this can be directly realized by evaluating the image at a number of focus positions and by an according extension of the merit function. Alternatively, the maximization of process windows has been demonstrated for a number of examples. This can be considered as the most rigorous approach and is especially well suited when co-optimizing for a number of different feature types that are to exhibit the largest possible common process regime. It should be emphasized that any of these criteria can be easily tailored to tool or process-related requirements.

As emphasized before, our approach not only accounts for the non-convexity of the search space but in fact exploits the multimodality to concurrently explore several promising search space regions. This procedure not only leads to a robust convergence behavior, whose final results are highly reproducible, but also to a final set of similarly fit solutions. This enables the decision makers to review the final solutions and to select a candidate best matching their preferences. To further develop this option, we have reformulated the source/mask optimization task as a multi-objective problem. There, instead of optimizing for an aggregation of the incommensurable sub-objectives, a strict partial ordering is applied, and the best performing candidate solutions can be considered as ideal compromises between the criteria. Out of the final set of these compromise solutions, the decision makers may select the solution that most closely meets their preference. The main rationale of this routine can be summarized as follows:

An a priori formal articulation of certain preferences is often difficult or even impossible. For example, in a weighted sum approach the decision maker has to pre-determine the parameters by which the sub-objectives are weighted. Often, however, the consequence of a specific weight set is not predictable but introduces an unintended bias; especially since the impact of a weight configuration may strongly change as the optimization progresses. Moreover, as shown in this thesis, scalar replacement strategies based on agglomeration functions omit (non-convex) compromise solutions that potentially perfectly reflect the preference profile of the decision maker. In this work, we have successfully applied a state-of-the-art multi-objective genetic algorithm to source/mask optimization problems, and we have clearly shown its ability to yield various adequate solutions. As a possible future extension to this approach, a semi-automatic post-optimization selection routine—for example, examining the sensitivity of results (as proposed by Barenbaum, 2010)—could be applied in order to further reduce the results presented to the decision maker. Moreover, recently devised indicator-based multi-objective evolutionary approaches may constitute a viable alternative to the direct non-dominated sorting approaches studied here.

Applied to the case of contact holes, both the single and the multi-objective evolutionary algorithm approaches were able to generate innovative, highly well-performing source/mask solutions. Specifically, all solutions exhibit both large exposure latitudes and extremely wide depths of focus without any risk of sidelobe printing. The examples also demonstrate that the approach can be directly applied to the co-optimization of different features and feature set-ups, for which a common ideal source was found.

We have also shown how the approach devised here can be extended to the domain of the projector in a straightforward manner, giving rise to a source/mask/projector optimization (SMPO) procedure. The motivation for this technique is to compensate for aberration-like mask-induced phase effects, which lead to a pronounced asymmetry of process windows, variably affecting different features. With the continuing downscaling and the narrowing process margins, these effects become more critical. On the other hand, pattern multiplication techniques drastically increase overlay demands and hence also the requirements for large dose and focus budgets. It is thus even more important to counteract mask phase artifacts. For that purpose, we and others have proposed to exploit the aberration control present in some of the state-of-the-art scanners. Owing to the flexibility of the optimization technique and framework developed in this work, we were the first to present a comprehensive SMPO approach and a number of well performing solutions demonstrating its feasibility. Particularly, we have explicitly demonstrated the need to rigorously co-optimize all three components: the source, the mask and the projector. In addition, such an optimization requires the use of advanced model options: The mask, for example, is to be computed using a full EMF solver, and spectra at different incidence angles need to be accounted for. We have demonstrated that these simulation options can be effortlessly integrated into our optimization approach.

Alternatively, the mask-induced EMF effect can be compensated by direct modifications of the mask. We have demonstrated this approach by a novel, rigorous optimization of the topology and material properties of the absorber stack. In addition to process-window-oriented optimization criteria, similar to those used in the source/mask optimization case, figures of merit specifically aimed at the mask performance—such as the mask-error-enhancement factor, and material constraints—have been introduced. The inclusion of further criteria, including a polarization sensitivity measure, have been shown. Like in the classical SMO case, the performance of the mask strongly depends on the illumination settings, which therefore also in this case have been co-optimized. Here, only a simplistic parametric representation is used. Future versions should replace this representation with the more flexible polar-coordinate-representation, which more stringently accounts for the illuminator flexibility of today's scanners. In order to render the absorber stack appropriate for a range of different feature set-ups, both dense and isolated lines/spaces have been simultaneously accounted for. To this end, a parametric assist feature placement routine was included into the optimization procedure, further increasing the degree of complexity. Yet the results conceivably demonstrate that this does not present an intricate problem for our optimization routine. Instead, numerous well-performing solutions exhibiting low sensitivities against mask size and polarization deviations, and substantially increased process windows were obtained. The same procedure was also applied to different contact hole set-ups, where similarly process budget improvements were achieved. Again, these improvements can be attributed to mainly three factors: (1) an increase of process latitudes of the individual features, (2) a decrease of an iso-

dense bias and (3) a reduction of feature-dependent process window asymmetries. In other words, the net outcome can be considered a joint effect of traditional resolution enhancement techniques (RET) and optical proximity correction (OPC) measures and the absorber adaptation. It is thus evident that all of these interdependent factors need to be accounted for simultaneously and require a irrefragably mutual optimization like the one proposed here.

The optimization strategy was further enhanced by replacing the evolutionary algorithm with a genetic search approach. There, a genetic algorithm serves as a global optimizer, and selected individuals are refined using a local search approach. For that purpose, a comprehensive hybridization framework has been developed in the scope of this work, and its performance was tested on a large number of benchmark functions. The results suggest an accelerated convergence and an increased reproducibility of results compared to a plain genetic algorithm (GA). This is noteworthy since the GA was equipped with a niching technique, which attempts to concurrently maintain multiple local optima. It can hence be speculated that the local search is capable of increasing the precision of the niche-specific solutions. The tests also revealed a strong dependence of the performance of the proposed hybrid approach, called memetic algorithm (MA), on its configuration parameters. Especially the interval between local searches and the mechanism by which solutions are selected for local improvements highly impact both the convergence speed and the precision. Our results confirm observations by others that show that too intensive a local search is often counterproductive in that it impedes the EA-driven exploration. In addition to a study of different MA configurations, we have also compared two different local search approaches: a state-of-the-art sequential quadratic program (SQP) algorithm and Nelder and Mead's direct search technique. We have shown that the faster convergence of the SQP method in comparison to the direct search does not necessarily lead to an improvement of the overall performance of the MA. Specifically, in order for the benchmark tests to more closely resemble the lithographic optimization problems aimed at, we have not provided any derivatives but left their numerical evaluation to the SQP method. This, of course, led to a significant increase of the number of function evaluations, which in turn degraded the efficiency of the MA. The Nelder-Mead (NM) technique, on the other hand, exploits only zero order information, maintaining a relatively small number of function calls, albeit at the price of a less intense search. This, however, did not present a considerable disadvantage for the MA. Nonetheless, if derivatives are available or computationally inexpensive, the SQP algorithm or related methods can be expected to be beneficial. Similarly, we have not conducted tests on optimization problems with non-linear constraints, for which the SQP method can be assumed to show an even greater advantage. But the formulation of constraints in the scope of MAs is far from being straightforward since it has to be posed in a way that is suited for both the numerical optimization routine and the GA. Neither have we considered adaptive approaches, proposed by other authors, in which the selection of local search approaches is performed heuristically or as a result of a co-evolution. Despite the well-founded rationale behind this approach, the adaptation of the search to local properties of the search space, the results presented by the authors are encouraging only for a small range of problems. Particularly, such a technique involves a considerable parameter tuning effort comparable to the thorough preselection of a suited local search method, thus countermanding the original goal of a "self-fulfilling prophecy." A promising direction also accounting for the local characteristics of the search function would be a tighter coupling of the local search and the employed niching method, for example, by conducting a local tournament as a Baldwinian competition in which not the best performing individual but the best performing candidate after the local search wins.

The benchmark tests showed that the local search operation is a significant contributor to the number of function evaluations of the MA. In order to reduce it, a function approximation technique using an artificial neural network (ANN) has been devised. To avoid overfitting and to adapt the approximated search space more closely to the dynamics presented by the GA-sampled landscape, network training is performed with the populations occurring between two local search generations. Different network topologies including one and two hidden layers have been investigated. Both layouts are suited to employ ANNs as universal function approximators. The single hidden layer showed a slightly attenuated sensitivity against variations of the number of nodes. Different learning algorithms were compared. A marginal advantage of a back-propagation algorithm with adaptive learning rates over a standard back-propagation routine was observed. The suitability of the approximated function was verified through

a number of benchmark experiments. The number of function approximations, especially their reduction in comparison to the non-approximated MA, strongly depends on the problem. For example, in the case of a highly multimodal test function (Schwefel's function), a significantly faster convergence (fitness versus function evaluations) could be achieved, whereas in the case of a function featuring a flat plateau with flat gradients, the performance of both algorithms was identical. The scaling of the input parameters was identified as a critical setting for the ANN approximation. Specifically, for the local search to function properly it should be chosen such that it leads to a wide spread of parameter values.

The memetic algorithm both with and without function approximation was applied to the optimization of the mask absorber stack for line/space patterns, in both cases showing an improvement of the maximum fitness of solution candidates compared with the plain genetic algorithm. The degree of improvement, however, is less significant than for the numerical benchmark tests. The highly fit results obtained with the plain GA and the large number of function evaluations required by the local search steps would render this example virtually unsuited for an application of the memetic algorithm. The use of function approximations and the associated drastic decrease of fitness evaluations, on the other hand, transforms the MA into a viable alternative to the stand-alone evolutionary algorithm. It is noteworthy that under function approximation neither the precision nor the quality of the final results suffer from a degradation. Contrarily, the approximated MA slightly outperforms both the GA and the fully evaluated MA. Further studies on adequate MA configuration parameters present an avenue for further investigations on performance increase potentials.

In the extreme-ultraviolet (EUV) case, because of the tight energy budget and the telecentricity errors due to the non-nominal incidence onto the mask, the absorber stack properties are even more critical. We have thus applied the previous optimization approach to the acquisition of ideal EUV mask stacks and illumination conditions for line/space and contact hole set-ups. Similarly as in the optical case, the goal was to optimize the common process window overlap of different pitches. Additionally, the joint process for different line widths was maximized, implying absorber candidates that exhibit a near-linear response to mask CD variations. In all examples, a substantial process window increase was achieved. Moreover, owing to the multi-objective optimization approach, a number of equally well performing solutions with different material and topology properties were obtained. It is interesting to note that alternating phase-shift masks with highly reflective absorbers (reflectance value of up to fifty percent) were among the superior solutions. As demonstrated, however, their strong sensitivity against small thickness variations does not render this mask type a viable option. This fact is implicitly accounted for by the evolutionary algorithm by maintaining only a small fraction of respective solutions. In addition to the attenuated, also alternating phase-shifting masks were considered in this study. There, the absorber stack parameters included the number of etched layers and the ratio between the shifter and the unshifted line. That study specifically emphasized the horizontal/vertical bias, which is a common problem under the non-telecentric EUV conditions. Solutions with minimum imbalances and hence showing large process window overlaps were obtained. As purported before, both the optimization approach and the accompanying infrastructure greatly facilitated the realization of such a multitude of surveys. Future applications may extend the approach to additional evaluation criteria, such as the multilayer stack or defects, or may be aimed at a full-fledged 3-D OPC procedure, in which a rigorous SMO routine is complemented by an absorber property co-optimization.

This high degree of flexibility of the optimization scheme developed here was once more demonstrated through a last application example, in which the anti-reflective performance of a bottom anti-reflective coating (BARC) was optimized. Aimed at double patterning applications, the goal was to obtain a BARC stack that maintains its anti-reflective properties under the first exposure. In order to account for the topographic modifications resulted from this first exposure, rigorous EMF simulations were included for the first time in this type of optimization. The optimization scheme facilitates the exploration of a multitude of different stacks, including single and bi-layer configurations. The multi-objective optimization approach allowed to identify a large variety of different well functioning stacks, from which in a post-optimization selection procedure the most promising compromises between manufacturability and performance could be identified. A number of alternative problem formulations were demonstrated and numerous innovative solutions were presented.

The previous examples clearly demonstrate the feasibility and the potentials of the proposed approach to rigorously optimize various aspects of the lithographic process. In addition to the multitude of evaluation criteria and constraints, virtually any other figure of merit can be included in a straightforward manner. For example, in the source/mask optimization case, a more rigorous treatment of the photoresist behavior would constitute a valuable alternative to the aerial-image-based evaluation applied in this work. In order to arrive at feasible computation times, compact or behavioral resist models can be expected to be best suited. As shown, an exact EMF treatment of both mask and wafer diffraction effects becomes increasingly important. Though directly integrable with our approach, full EMF simulations—especially when accounting for different incidence angles—drastically increase the required computation times. In this work, we have attempted to counteract this necessary performance degradation by an application of speed-up techniques such as domain decompositions. To further improve the utility of the proposed approach, more aggressive approximative techniques, including compact models, might be required.

Other accelerative techniques could be further developed. For example, the neural network approach used for function approximations in the local search stage of the memetic algorithm could be extended to the entire evolutionary algorithm. This would of course also require an upgrade of the model control infrastructure in order to ensure that the search is conducted on a viable estimate representation of the actual search space. Other speed-up potentials are related to restricting the search space by additional problem-specific knowledge. Along that line, both source and mask candidate solutions could be constrained according to their spatial frequency response. For instance, given the diffraction spectrum of a mask candidate, the number of potentially useful source solutions drastically diminishes.

The optimization cases presented in this work are tailored to lithography simulation problems and do not specifically account for the large-area requirements that are common in the majority of EDA applications. More specifically, in this work a mesoscopic perspective has been taken, focusing on the performance of critical features and feature configurations. To extend this to a full-chip-oriented approach, a number of considerations would need to be added: including a spatial-frequency-aware pattern recognition engine discovering identical feature configurations to re-use evaluation results and a decomposition and stitching framework by which individual small-area computations are recombined to form the large-area solution. Moreover, the flexibility of our optimization approach would also allow to employ coarser, faster approximations such as the Hopkins instead of the Abbe imaging approach. A more rigorous exploitation of hierarchical fitness evaluations by which abysmal candidates are discarded without conducting an exhaustive simulation could also prove beneficial. A potential avenue for further research would also include a machine-learning approach in which rules are constantly yielded from rigorously evaluated training patterns. Instead of directly optimizing the mask, the layout is modified through these rules. Such a technique would be comparable to traditional rule-based OPC approaches but where the rules are generated automatically.

Finally, a dedicated multi-objective memetic algorithm, for which a number of approaches have been proposed over the past years, may be a valuable addition in that it may speed up the multi-objective optimization approaches.

A Derivation of the Scalar Wave Equations

In differential form and MKS (meter-kilogram-second) notation, Maxwell's equations are:¹

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t} \quad \text{Faraday's law} \quad (\text{A.1a})$$

$$\nabla \times \mathbf{H} = \mathbf{J} + \frac{\partial \mathbf{D}}{\partial t} \quad \text{Ampère's law with Maxwell's correction} \quad (\text{A.1b})$$

$$\nabla \cdot \mathbf{D} = \rho \quad \text{Gauss's law} \quad (\text{A.1c})$$

$$\nabla \cdot \mathbf{B} = 0 \quad \text{Gauss's law for magnetism,} \quad (\text{A.1d})$$

Although describing the principle of the relation between the electric displacement field \mathbf{D} , the electric field \mathbf{E} , the magnetizing field \mathbf{H} , the magnetic induction \mathbf{B} and the current density \mathbf{J} , Equations (A.1a)–(A.1d) are not sufficient to actually compute the impact of currents or charges on the field, unless supplemented with material properties.² For the majority of lithography applications, materials can be assumed to be linear, isotropic, time-invariant and inhomogeneous. Also no free charges occur. The constitutive relations can hence be stated as follows:

$$\mathbf{D} = \varepsilon \mathbf{E} \quad (\text{A.2a})$$

$$\mathbf{H} = \mathbf{B}/\mu \quad (\text{A.2b})$$

$$\mathbf{J} = \sigma \mathbf{E} \quad \text{Ohm's law,} \quad (\text{A.2c})$$

where ε is the electric permittivity, μ is the magnetic permeability and σ is the conductivity of the medium. In general, these properties are not constants but functions (evident, for example, for inhomogeneous materials).³ Substituting (A.2a) and (A.2b) into the curl equations (A.1a) and (A.1b), we have that

$$\nabla \times \mathbf{E} = -\mu \frac{\partial \mathbf{H}}{\partial t} \quad (\text{A.3a})$$

$$\nabla \times \mathbf{H} = \sigma \mathbf{E} + \varepsilon \frac{\partial \mathbf{E}}{\partial t} \quad (\text{A.3b})$$

Application of the curl operator to (A.3a) leads to:

$$\begin{aligned} \nabla \times (\nabla \times \mathbf{E}) &= -\mu \frac{\partial}{\partial t} (\nabla \times \mathbf{H}) \\ &= -\mu \frac{\partial}{\partial t} (\sigma \mathbf{E} + \varepsilon \frac{\partial \mathbf{E}}{\partial t}) \\ &= -\mu \sigma \frac{\partial \mathbf{E}}{\partial t} - \mu \varepsilon \frac{\partial^2 \mathbf{E}}{\partial t^2} \\ &= \nabla (\nabla \cdot \mathbf{E}) - \nabla^2 \mathbf{E} \quad (\text{using } \nabla \times (\nabla \times \mathbf{v}) = \nabla (\nabla \cdot \mathbf{v}) - \nabla^2 \mathbf{v}) \end{aligned} \quad (\text{A.4})$$

For charge-free materials ($\nabla \cdot \mathbf{E} = 0$) and non-absorbing media ($\sigma = 0$), one hence gets:⁴

$$\nabla^2 \mathbf{E} = \mu \varepsilon \frac{\partial^2 \mathbf{E}}{\partial t^2} \quad (\text{A.5})$$

¹(for a derivation (there, in Gaussian units) see, e.g., Jackson, 1962, chap. 6, pp. 177–179)

²See, for example, Born and Wolf (1999, chap. 1, pp. 2–4).

³The next derivation steps essentially follow Born and Wolf (1999, chap. 1, pp. 11–15) and Goodman (1996, chap. 3, pp. 36 and 37).

⁴also compare Jackson (1962, chap. 7, pp. 202).

Similarly for the magnetic field (where $\nabla \cdot \mathbf{H} = 0$):

$$\nabla^2 \mathbf{H} = \mu \varepsilon \frac{\partial^2}{\partial t^2} \mathbf{H} \quad (\text{A.6})$$

Resulting in the wave equations:

$$\nabla^2 \mathbf{E} = \frac{n^2}{c_0^2} \frac{\partial^2}{\partial t^2} \mathbf{E} \quad (\text{A.7})$$

$$\nabla^2 \mathbf{H} = \frac{n^2}{c_0^2} \frac{\partial^2}{\partial t^2} \mathbf{H}, \quad (\text{A.8})$$

where n is refractive index of the medium, defined as

$$n = \sqrt{\frac{\mu \varepsilon}{\mu_0 \varepsilon_0}} = \frac{c_0}{v_p} \quad (\text{A.9})$$

where c_0 the speed of light in vacuum (with $c_0 = \sqrt{\mu_0 \varepsilon_0}$) and ε_0 is the permittivity in vacuum and μ_0 the permeability in vacuum. ($v = 1/\sqrt{\mu \varepsilon}$ introduces the concept of velocity of a wave.)

In a dielectric, linear, isotropic, homogeneous, and nondispersive medium, the vector components of the former wave equations behave identically, yielding the following scalar wave equation:⁵

$$\nabla^2 u = \frac{1}{v_p^2} \frac{\partial^2 u}{\partial t^2}. \quad (\text{A.10})$$

In the case of a monochromatic field, the scalar field can be explicitly given by the harmonic motion equation:⁶

$$u(\mathbf{r}, t) = A(\mathbf{r}) \cos(\phi(\mathbf{r}) + \omega t), \quad (\text{A.11})$$

where A is the amplitude, ϕ is the phase, and ω is the angular frequency.

Alternatively, (A.11) can be formulated as follows:

$$\begin{aligned} u(\mathbf{r}, t) &= \Re\{A(\mathbf{r}) \exp[-j(\phi(\mathbf{r}) + \omega t)]\} \\ &= \Re\{\underbrace{A(\mathbf{r}) \exp[-j\phi(\mathbf{r})]}_{U(\mathbf{r}) \text{ (phasor)}} \exp[-j\omega t]\} \\ &= \Re\{U(\mathbf{r}) \exp(-j\omega t)\} \end{aligned} \quad (\text{A.12})$$

With that, a separation of the spatial and the temporal parts of u is obtained. The differential terms of the wave equation (A.10) can now be stated as

$$\nabla^2 u = \Re\{\exp(-j\omega t) \nabla^2 U\} \quad (\text{A.13})$$

and

$$\frac{\partial^2 u}{\partial t^2} = -\omega^2 \Re\{U \exp(-j\omega t)\}. \quad (\text{A.14})$$

Substitution in (A.10) and using the wavenumber notation ($k = \omega/c$) yields the time-independent Helmholtz equation

$$\begin{aligned} \Re\{\exp(-j\omega t) \nabla^2 U\} &= -\frac{\omega^2}{c^2} \Re\{U \exp(-j\omega t)\} \\ \nabla^2 U &= -\frac{\omega^2}{c^2} U \\ (\nabla^2 + k^2)U &= 0. \end{aligned} \quad (\text{A.15})$$

⁵Of course, the following equation has no unique but many solutions, including straightforward ones that lead to *plane* and *spherical waves*.

⁶Compare Goodman (1996, chap. 3, pp. 38 and 39) for the following steps.

B Diffraction

The mathematical description of diffraction in imaging systems such as microscopes or lithographic projection systems are often based on either the Fresnel-Kirchoff or the Rayleigh-Sommerfeld formulae.¹ Both theorems are based on an integral form introduced by Helmholtz and Kirchhoff, which is based on Green's theorem. The basic idea is depicted in Figure B.1: The disturbance at an observation point P is expressed in terms of the surface (S) encompassing a volume (V), which the point is situated in.




Figure B.1: Schematic of initial point leading Helmholtz-Kirchhoff integral theorem.

In three dimensions, Green's theorem can be stated as follows:

$$\iiint_V (U \nabla^2 \hat{U} - \hat{U} \nabla^2 U) dV = \iint_S \left(U \frac{\partial \hat{U}}{\partial \mathbf{n}} - \hat{U} \frac{\partial U}{\partial \mathbf{n}} \right) dS, \quad (\text{B.1})$$

where, U and \hat{U} are functions with single-valued and continuous first and second order derivatives, V and S are the volume and surface, respectively, as shown in Figure B.1. The derivative in the outward normal direction on S is denoted by $\partial/\partial \mathbf{n}$. An auxiliary function \hat{U} has to be selected, which, in addition to the former requirements regarding continuity and differentiability, has to fulfill the Helmholtz wave equation. This is of course also true for the sought function, so that we have:

$$\begin{aligned} \nabla^2 U &= -k^2 U \quad \text{and} \\ \nabla^2 \hat{U} &= -k^2 \hat{U}. \end{aligned} \quad (\text{B.2})$$

Using this property and substituting into Green's equation (B.1), yields:

$$\iint_S \left(U \frac{\partial \hat{U}}{\partial \mathbf{n}} - \hat{U} \frac{\partial U}{\partial \mathbf{n}} \right) dS = 0. \quad (\text{B.3})$$

Kirchhoff proposes the following Green's function:

$$\begin{aligned} \hat{U}(P') &= \frac{\exp(jk|P' - P|)}{|P' - P|} \\ \hat{U}(P') &= \frac{\exp(jkr)}{r} \quad \text{with } r = |P' - P|, \end{aligned} \quad (\text{B.4})$$

¹The derivations in this appendix follow Born and Wolf (1999, chap. 8) and Goodman (1996, pp.39–50).

where r is the distance between the observation point and the argument of \hat{U} . \hat{U} exhibits a singularity at the observation point, i.e., for $P' = P$. This discontinuity can be resolved by introducing an infinitesimal volume V_ε , which surrounds P , and excluding this volume from the disturbance calculation. Thus also a new surface consisting of the original surface and a surface S_ε that bounds the epsilon volume has to be used. Equation (B.3) becomes:

$$\iint_S \left(U \frac{\partial \hat{U}}{\partial \mathbf{n}} - \hat{U} \frac{\partial U}{\partial \mathbf{n}} \right) dS = - \iint_{S_\varepsilon} \left(U \frac{\partial \hat{U}}{\partial \mathbf{n}} - \hat{U} \frac{\partial U}{\partial \mathbf{n}} \right) dS_\varepsilon. \quad (\text{B.5})$$

For points P' situated on the surface S_ε , which have a distance from the observation point P of ε , we obtain

$$\begin{aligned} \hat{U}(P') &= \frac{e^{jk\varepsilon}}{\varepsilon} \\ \frac{\partial \hat{U}(P')}{\partial \mathbf{n}} &= \frac{e^{jk\varepsilon}}{\varepsilon} \left(\frac{1}{\varepsilon} - jk \right). \end{aligned}$$

Letting ε approach zero, we have

$$\iint_{S_\varepsilon} \left(U \frac{\partial \hat{U}}{\partial \mathbf{n}} - \hat{U} \frac{\partial U}{\partial \mathbf{n}} \right) dS_\varepsilon \xrightarrow{(\varepsilon \rightarrow 0)} 4\pi\varepsilon^2 \left(U \frac{e^{jk\varepsilon}}{\varepsilon} \left(\frac{1}{\varepsilon} - jk \right) \frac{e^{jk\varepsilon}}{\varepsilon} \frac{\partial U}{\partial \mathbf{n}} \right) \xrightarrow{(\varepsilon \rightarrow 0)} 4\pi U.$$

Finally, substituting in (B.5), yields:

$$U(P) = -\frac{1}{4\pi} \iint_S \left(U \frac{\partial \hat{U}}{\partial \mathbf{n}} - \hat{U} \frac{\partial U}{\partial \mathbf{n}} \right) dS \quad (\text{B.6})$$




Figure B.2: Fresnel-Kirchhoff diffraction principle.

In order to present a more concise mathematical formulation for the Huygens-Fresnel principle, Kirchhoff applied the integral theorem (B.6) to the situation depicted in Figure B.2. The goal is again to determine the disturbance at a point P located in the center of a sphere (S_2) with radius R . A wave originating from a point P_0 impinges on an opaque screen (S_1) with an aperture (Σ). The surface

surrounding the observation point is hence given by the sum of S_1 —not including the aperture— S_2 and the aperture Σ . Applying the integral theorem (B.6), leads to

$$U = \frac{1}{4\pi} \left(\iint_{S_1} \iint_{S_2} \iint_{\Sigma} \right) \left(U \frac{\partial \hat{U}}{\partial n} - \hat{U} \frac{\partial U}{\partial n} \right) dS. \quad (\text{B.7})$$

For a sphere S_2 with a large radius, its contribution to the disturbance at P vanishes. This can be deduced as follows: By noting that the derivative of \hat{U} approximates to $jk\hat{U}$ for large values of R and by expressing the disturbance due to S_2 in terms of the solid angle Ω that S_2 subtends at P , we obtain:

$$\iint_{S_2} \left(\frac{\partial U}{\partial \mathbf{n}} \frac{e^{jkR}}{R} - U \frac{e^{jkR}}{R} \left(jk - \frac{1}{R} \right) \right) dS \approx \iint_{S_2} \hat{U} \left(\frac{\partial U}{\partial \mathbf{n}} - U jk \right) dS \equiv \int_{\Omega} \hat{U} \left(\frac{\partial U}{\partial \mathbf{n}} - U jk \right) R^2 d\Omega. \quad (\text{B.8})$$

Applying the Sommerfeld radiation condition, that is,

$$R \left(\frac{\partial U}{\partial \mathbf{n}} - U jk \right) \xrightarrow{R \rightarrow \infty} 0, \quad (\text{B.9})$$

leads to the claimed extinction of an S_2 contribution. The integral hence reduces to:

$$U = \frac{1}{4\pi} \iint_{S_1 + \Sigma} \left(\hat{U} \frac{\partial U}{\partial n} - U \frac{\partial \hat{U}}{\partial n} \right) dS. \quad (\text{B.10})$$

In order to solve this equation, Kirchoff proposed the following boundary conditions:

1. Within the aperture, the field is undisturbed.
2. On the backside of the screen S_2 , which is facing toward P , both the field and its derivative are zero.

Applying these boundary conditions, assuming that the distances from the wave source to screen S_1 and from S_1 to P are a multitude of the wavelength (λ), and assuming a single spherical wave with amplitude A , we have the Fresnel-Kirchhoff diffraction equation:

$$U(P) = \frac{jA}{\lambda} \iint_{\Sigma} \frac{e^{jk(r_0+r_1)}}{r_0 r_1} \frac{\cos(\mathbf{n}, \mathbf{r}_1) - \cos(\mathbf{n}, \mathbf{r}_0)}{2} d\Sigma. \quad (\text{B.11})$$

By setting

$$U'(Q) = \frac{1}{j\lambda} \frac{A e^{jk r_0}}{r_0} \frac{\cos(\mathbf{n}, \mathbf{r}_1) - \cos(\mathbf{n}, \mathbf{r}_0)}{2}$$

(B.11) can be rewritten as,

$$U(P) = \iint_{\Sigma} U'(Q) \frac{e^{jk r_1}}{r_1} d\Sigma \quad (\text{B.12})$$

The latter formula can indeed be regarded as a mathematically correct version of the phenomenological Huygens-Fresnel principle: The field in P is a result of an infinite number of secondary source points, whose amplitudes and phases result from the original illumination source and the angles of incidence.

Using an alternative Green's function, which employs both the original source point P and the same point mirrored along the screen and with a 180 degree phase difference, that is,

$$\hat{U} = \frac{e^{jkr}}{r} - \frac{e^{-jkr'}}{r'}, \quad (\text{B.13})$$

eliminates the requirement for boundary conditions when applied to the problem depicted in Figure B.2, as it vanishes not only on the backside of the screen but also in the aperture. This solution, found by

Sommerfeld, resolves a number of inconsistencies that result from the Kirchhoff boundary conditions. The corresponding diffraction formula, termed Rayleigh-Sommerfeld, is as follows

$$U(P) = \frac{A}{j\lambda} \iint_{\Sigma} \frac{e^{jk(r_0+r_1)}}{r_0 r_1} \cos(\mathbf{n}, \mathbf{r}_1) d\Sigma. \quad (\text{B.14})$$

Given Cartesian coordinates, (η, ξ) in the aperture plane and (x, y) in the observation plane, the Huygens-Fresnel principle can be written as follows:

$$U(x, y) = \frac{z}{j\lambda} \iint_{\Sigma} U'(\eta, \xi) \frac{e^{jk r_1}}{r_1^2} d\eta d\xi, \quad (\text{B.15})$$

where the distance r_1 from a point on the aperture to the observation screen is given by:

$$r_1 = \sqrt{z^2 + (x - \eta)^2 + (y - \xi)^2}$$

Using a binomial expansion for r_1 and truncating it to the first two terms in the exponent and to only the first one in the denominator, we obtain the Fresnel approximation:

$$U(x, y) = \frac{e^{jkz}}{j\lambda z} \iint_{-\infty}^{\infty} U'(\eta, \xi) \exp\left(\frac{jk}{2z} ((x - \eta)^2 + (y - \xi)^2)\right) d\eta d\xi, \quad (\text{B.16})$$

The Fresnel approximation can be said to be valid if the Fresnel number is larger or equal to one, that is,

$$F = \frac{a^2}{L\lambda} \geq 1, \quad (\text{B.17})$$

where a is the size of the aperture and L is the distance from the aperture.

An additional simplification is presented by the Fraunhofer diffraction formula, which, apart from a phase correcting factor, is equivalent to a Fourier transform of the aperture.

$$U(x, y) = \frac{e^{jkz} e^{\frac{jk}{2z}(x^2+y^2)}}{j\lambda z} \iint_{-\infty}^{\infty} U'(\eta, \xi) e^{-j\frac{2\pi}{\lambda z}(\eta x + \xi y)} d\eta d\xi, \quad (\text{B.18})$$

which is valid for $F \ll 1$.

C Scalar Imaging

In this appendix, the scalar imaging equation for a partially coherent illumination source, which presents the foundation of all lithography imaging models, is briefly derived. It is based largely on Born and Wolf (1999, chap. 10) and Wong (2005, chap. 4).

C.1 Coherent Imaging

Assuming coherent illumination, the projection system acts as a linear filter. Given spatial frequencies (f, g) , the field transmitted by the object U_0 and a response function of the imaging system H , the corresponding field in the image plane is given by:

$$\mathcal{F}(U_1)(f, g) = \mathcal{F}(U_0)(f, g)\mathcal{F}(H)(f, g). \quad (\text{C.1})$$

The response function H is:

$$H(x_1 - x_0, y_1 - y_0) = \frac{A}{\lambda z_1} \iint_{-\infty}^{\infty} P(\eta, \xi) \exp\left(-i \frac{2\pi}{\lambda z_1} ((x_1 - x_0)\eta + (y_1 - y_0)\xi)\right) d\eta d\xi, \quad (\text{C.2})$$

where (x_0, y_0) is a point in the object plane and (x_1, y_1) is the point in the image plane at which the (x_0, y_0) contribution is to be evaluated. Here, x_0 and y_0 are transformed such that image and object coordinates correspond in the following way:

$$x_0 := Mx'_0 \quad \text{and} \quad y_0 := My'_0, \quad (\text{C.3})$$

where (x'_0, y'_0) are the Cartesian coordinates in the object space and M is the magnification factor. In (C.2), the distance of the image plane from the pupil plane is given by z_1 . The factor A is a constant amplitude scaling factor. P denotes the pupil function, and the coordinates (η, ξ) refer to the pupil plane. The pupil function is unity for (η, ξ) points within the pupil's aperture and zero otherwise.

The response function H (sometimes referred to as amplitude point-spread function) represents the Fraunhofer diffraction pattern of the pupil. Thus its Fourier transform is equal to the pupil function itself (scaled by the product of the constant amplitude, the wavelength and the distance from the image plane) if the following coordinate transformation is applied (Born and Wolf, 1999, sec. 9.5, p. 546):

$$\eta = -\lambda z_1 f \quad \text{and} \quad \xi = -\lambda z_1 g$$

In the case of circular apertures—as common to lithographic lens systems—the pupil function can be formally defined as follows:

$$P : (f, g) \mapsto \begin{cases} 1 & \sqrt{f^2 + g^2} \leq \frac{a}{\lambda z_1} \\ 0 & \text{otherwise} \end{cases},$$

where a is the radius of the aperture.

The lithographic illumination is neither completely coherent nor incoherent. Therefore the mutual intensity and the mutual degree of coherence caused by the illumination source and the impact on the propagated light have to be taken into account.




Figure C.1: Intensity distribution in dependence of mutual coherence of two points.

C.2 Mutual Intensity

The following relationships are defined for the situation depicted in Figure C.1, in which two pinholes P_1 and P_2 on a screen \mathcal{S}_1 are illuminated from an extended source and where the intensity on a second screen \mathcal{S}_2 is to be determined.

The *mutual coherence function* Γ is obtained as follows:¹

$$\Gamma : (P_1, P_2, \tau) \mapsto \langle U(P_1, t + \tau) U^*(P_2, t) \rangle.$$

The distance of the point of interest Q from P_1 and P_2 are s_1 and s_2 , respectively. The potentially resulted time difference of the traveling light is taken into account by τ .

Similarly, the *complex degree of coherence* γ is defined as:

$$\gamma : (P_1, P_2, \tau) \mapsto \frac{\Gamma(P_1, P_2, \tau)}{\sqrt{I_1} \sqrt{I_2}},$$

where I_1 and I_2 denote the independent intensities at P_1 and P_2 , respectively.

By neglecting the infinitesimal time difference resulted from the different path lengths s_1 and s_2 (viz. $(s_2 - s_1)/c$), we can set:

$$\begin{aligned} J(P_1, P_2) &:= \Gamma(P_1, P_2, 0) \text{ and} \\ j(P_1, P_2) &:= \gamma(P_1, P_2, 0). \end{aligned}$$

C.3 Imaging Under Illumination of an Extended Incoherent Source

Given an extended source Σ that is decomposed into n mutually incoherent and randomly phase-related elements $d\sigma_i$ ($i = 1, \dots, n$), with $U_i(P, t)$ being the field at point P induced by $d\sigma_i$ (see Figure C.2). We then have that

$$U(P, t) := \sum_{i=1}^n U_i(P, t)$$

for the total field at point P .

The mutual intensity for the extended source is hence:

$$J(P_1, P_2) = \sum_{i=1}^n \langle U_i(P_1, t) U_i^*(P_2, t) \rangle + \sum_{i=1}^n \sum_{j=1; j \neq i}^n \langle U_i(P_1, t) U_j^*(P_2, t) \rangle \quad (\text{C.4})$$

¹The operator $\langle \cdot \rangle$ denotes the ensemble average. In this particular case, it is defined as the time average, i.e.:

$$\langle f(t) \rangle \stackrel{T \rightarrow \infty}{=} \frac{1}{-T} \int_{-T/2}^{T/2} f(t) dt,$$

where T is the time period of observation.




Figure C.2: Mutual intensity of points remote from extended incoherent source. The source Σ is decomposed into mutually incoherent elements $d\sigma_i$ with center σ . The mutual intensity of points P_1 and P_2 on a remote screen \mathcal{S}_1 is evaluated.

Because of the mutual incoherence of the source elements, the terms of the double sum on the right of (C.4) are equal to zero. The remaining summands can be expressed in terms of the amplitude emitted from the individual source element:

$$U_i(P) = A_i \left(t - \frac{|P - \sigma_i|}{c} \right) \frac{e^{-i\omega \left(t - \frac{|P - \sigma_i|}{c} \right)}}{|P - \sigma_i|} \quad \forall P \in \{P_1, P_2\},$$

where $A_i(t)$ is the amplitude of source element $d\sigma_i$, ω is the average angular frequency, and $|P - \sigma_i|$ is the distance from $d\sigma_i$ to P_1 and P_2 , respectively. We again disregard the time difference resulting from the path difference $|P_1 - \sigma_i| - |P_2 - \sigma_i|$ in the argument of A_i^* and obtain:

$$J(P_1, P_2) = \sum_{i=1}^n \langle A_i(t) A_i^*(t) \rangle \frac{e^{-i\omega \left(\frac{|P_2 - \sigma_i| - |P_1 - \sigma_i|}{c} \right)}}{|P_2 - \sigma_i| - |P_1 - \sigma_i|} \quad (\text{C.5})$$

Although, as we shall see in the next paragraph, numerical solutions indeed make use of a discretized version of the source, in reality, it has to be treated continuously. We define the intensity per unit area of the source as:

$$I(\sigma_i) d\sigma_i = \langle A_i(t) A_i^*(t) \rangle$$

and hence have the following form of the continuous mutual intensity:

$$J(P_1, P_2) = \iint_{\Sigma} I(\sigma) \frac{e^{-ik(|P_2 - \sigma_i| - |P_1 - \sigma_i|)}}{|P_2 - \sigma_i| - |P_1 - \sigma_i|} d\sigma. \quad (\text{C.6})$$

Likewise, the complex degree of coherence can be defined as:

$$j(P_1, P_2) = \frac{1}{\sqrt{I_1} \sqrt{I_2}} \iint_{\Sigma} I(\sigma) \frac{e^{ik(R_1 - R_2)}}{R_1 R_2} d\sigma. \quad (\text{C.7})$$

Equation (C.7) is sometimes referred to as *van Cittert-Zernike theorem*, which can be derived from it and which relates the complex degree of coherence to a Fresnel-Kirchhoff diffraction formula [for details, see for example (Hecht, 2002, p. 572) or (Born and Wolf, 1999, p. 574)].

C.4 Hopkins' Formula

Former equations [(C.6) and (C.7)] assume a homogeneous medium between the source and the pinholes. In order to generalize the theorem to heterogeneous media, the fixed transmission terms $\frac{e^{-i\omega(t-R/c)}}{R}$ (where R is the respective path length from the source) are to be replaced by a general transmission function $H(\sigma_i, P, \omega)$. If we also set:

$$U(\sigma, P) := \sqrt{I(\sigma)} H(\sigma, P, \omega) \quad \forall P \in \{P_1, P_2\},$$

(C.6) and (C.7) become

$$J(P_1, P_2) = \iint_{\Sigma} U(\sigma, P_1) U(\sigma, P_2) d\sigma \quad (\text{C.8})$$

and

$$j(P_1, P_2) = \frac{1}{\sqrt{I_1} \sqrt{I_2}} \iint_{\Sigma} U(\sigma, P_1) U(\sigma, P_2) d\sigma \quad (\text{C.9})$$

respectively. The former equations were first proposed by Hopkins (1951) and are therefore often termed *Hopkins' formulae*.

C.5 Propagation of Mutual Intensity




Figure C.3: Propagated mutual intensity of points Q_1, Q_2 under illumination of P_1 and P_2 from extended incoherent source Σ .

The next considerations are concerned with the propagation of the mutual intensity of pinholes (as depicted in Figure C.1) onto a second screen. More precisely, as shown in Figure C.3, given the mutual intensity of a pair of points (P_1, P_2) on a screen (S_1) illuminated from an incoherent extended source (Σ), the mutual intensity of point pairs Q_1, Q_2 on a second remote screen (S_2) can be determined by:

$$J_1(Q_1, Q_2) = \iint_{S_1} \iint_{S_1} J_0(P_1, P_2) H(P_1, Q_1) H^*(P_2, Q_2) dP_1 dP_2. \quad (\text{C.10})$$

Equation (C.10) can be derived by noting that the Huygens-Fresnel principle (generalized for a transmission function H from the first to the second screen) yields

$$U(\sigma, Q) = \iint_{S_1} U(\sigma, P) H(P, Q) dP \quad \forall (P, Q) \in \{(P_1, Q_1), (P_2, Q_2)\},$$

and by accordingly substituting $U(\sigma, Q_1)U^*(\sigma, Q_2)$ in $J_1(Q_1, Q_2)$ [cf. (C.8)]. Finally, by exchanging the integration order, the integral over Σ yields $J(P_1, P_2)$ and we obtain (C.10).²

The equations presented here constitute a generalization of the mutual intensity propagation equation developed by Zernike (1938) in that they incorporate an arbitrary transmission function (H) to account for the optical system between screens \mathcal{S}_1 and \mathcal{S}_2 .

C.6 Köhler Illumination

As mentioned before, the illumination system employed in lithography is adopted from a technique frequently used in microscopy and invented by Köhler. The idea is to reduce irregularities in the illumination by utilizing an image of it instead.

Let the field in points P_1 and P_2 in the object plane be given by:

$$U(\sigma, P_i) = Ae^{i\phi_i} \quad i = 1, 2, \quad (\text{C.11})$$

where A is the source amplitude and ϕ_i are the phases at the respective point $P_i = (x_i, y_i)$:

$$\phi_i = k(s_x x_i + s_y y_i),$$

with k the mean wavenumber and (s_x, s_y) the first components of the unit wave vector. With the phase difference:

$$\phi_1 - \phi_2 = k(s_x(x_1 - x_2) - s_y(y_1 - y_2))$$

we obtain

$$j(P_1, P_2) = \frac{1}{\sqrt{I(P_1)}\sqrt{I(P_2)}} \iint_{\Sigma} A^2 e^{-ik(\phi_1 - \phi_2)} d\sigma \quad (\text{C.12})$$

with

$$I(P_i) = \iint_{\Sigma} A^2 d\sigma.$$

C.7 Partially Coherent Imaging

Let $P_1 = (x_0, y_0)$ and $P_2 = (x'_0, y'_0)$ two points in the object plane with a mutual intensity $J_O(P_1, P_2) = J_O((x_0, y_0), (x'_0, y'_0))$. Moreover, we apply the same coordinate transformation as in (C.3) in order for the object points to have the same coordinates as the image points. According to (C.10) and applying an isoplanatic approximation, the mutual intensity at two points $Q_1 = (x_1, y_1)$ and $Q_2 = (x'_1, y'_1)$ in the image plane is given by:

$$J_I((x_1, y_1), (x'_1, y'_1)) = \iiint_{-\infty}^{\infty} J_O((x_0, y_0), (x'_0, y'_0)) H(x_1 - x_0, y_1 - y_0) H^*(x'_1 - x'_0, y'_1 - y'_0) dx_0 dy_0 dx'_0 dy'_0. \quad (\text{C.13})$$

To determine the mutual intensity emergent from the objective, let us consider the disturbance at a point P due to a source point σ under Köhler illumination ($\bar{U}(\sigma, P)$). Further, let $O(P)$ be the object transmission function, for example, of the lithographic photomask. The disturbance of the source point is thus given by the product of the object function and the original field:

$$U(\sigma, P) = \bar{U}(\sigma, P)O(P). \quad (\text{C.14})$$

²Further details are presented by Born and Wolf (1999, p. 581).

The mutual intensities of two points P_1, P_2 both from the condenser J_C and from the object J_O is given by Hopkins' formula (C.8):³

$$J_C(P_1, P_2) = \int_{\Sigma} \bar{U}(\sigma, P_1) \bar{U}^*(\sigma, P_2) dS \quad (\text{C.15})$$

$$J_O(P_1, P_2) = \int_{\Sigma} U(\sigma, P_1) U^*(\sigma, P_2) dS. \quad (\text{C.16})$$

Substituting (C.14) in (C.16) and using (C.15), we can express the mutual intensity emitting from the object in terms of the object function O and the mutual intensity of the points from the condenser:

$$J_O(P_1, P_2) = O(P_1) O^*(P_2) J_C(P_1, P_2). \quad (\text{C.17})$$

Under certain illumination conditions including Köhler illumination, the mutual intensity is dependent only on the difference rather than on the absolute coordinates of two points. Using this simplification and substituting (C.17) in (C.13), we have that for the intensity of (x_1, y_1) on the image plane:

$$\begin{aligned} I(x_1, y_1) &= J_I((x_1, y_1), (x_1, y_1)) \\ &= \iiint \int_{-\infty}^{\infty} J_C(x_0 - x'_0, y_0 - y'_0) O(x_0, y_0) O^*(x'_0, y'_0) \\ &\quad H(x_1 - x_0, y_1 - y_0) H^*(x_1 - x'_0, y_1 - y'_0) dx_0 dy_0 dx'_0 dy'_0. \end{aligned} \quad (\text{C.18})$$

Representing J_C , H and O by their Fourier integrals yields

$$\begin{aligned} I(x_1, y_1) &= \iiint \iiint \int_{-\infty}^{\infty} \mathcal{F}(J_C)(f, g) \mathcal{F}(O)(f', g') \mathcal{F}(O)^*(f'', g'') \\ &\quad \mathcal{F}(H)(f + f', g + g') \mathcal{F}(H)^*(f + f'', g + g'') \\ &\quad \exp(-i2\pi((f' - f'')x_1 + (g' - g'')y_1)) df dg df' dg' df'' dg''. \end{aligned} \quad (\text{C.19})$$

³Here the same coordinate convention as in (C.3) is assumed.

D Transfer Matrix Method

The transfer matrix method is a common numerical technique to evaluate transmittance and reflectance values in stratified media. The derivation here is adopted from Hecht (2002, chap. 9, pp. 426–428).

transfer matrix method (TMM) makes use a TE/TM decomposition of the wave. Figure D.1 illustrates the TE case for a stack consisting of only one film. There, two interfaces have to be taken into account: Interface I is between the medium from where the light is incident and the film, interface II constitutes the boundary between the film and the substrate.




Figure D.1: Thin-film layer consisting of one film, the medium of incidence and the substrate [each of which may have a different refractive index (n)]. The electric and magnetic field components are illustrated for the TE case; that is, the magnetic field is in the plane of incidence, the electric field points into the illustration plane. The \mathbf{k} vectors indicate the propagation direction of the waves. The corresponding incidence, reflection and transmission angles are denoted by θ . The film has a thickness of Δ_1 .

The tangential components of the electric and the magnetic field are continuous along the interfaces. For the first interface (I), this yields:

$$E_I := E_i + E_r = E_{t'} + E_{r'}. \quad (\text{D.1})$$

For non-magnetic media, the relation between the electric and the magnetic can be stated as follows:

$$\mathbf{H} = \sqrt{\frac{\epsilon_0}{\mu_0}} n \mathbf{k} \times \mathbf{E},$$

where the refractive index n is defined as

$$n = \sqrt{\frac{\mu\epsilon}{\mu_0\epsilon_0}}.$$

The magnetic field component can hence be expressed in terms of the electric field:

$$\begin{aligned} H_I &:= \sqrt{\frac{\epsilon_0}{\mu_0}} (E_i - E_r) n_0 \cos \theta_i \\ &= \sqrt{\frac{\epsilon_0}{\mu_0}} (E_{t'} - E_{r''}) n_1 \cos \theta_1. \end{aligned} \quad (\text{D.2})$$

Similar relations result at the second interface (II):

$$E_{II} := E_{i'} + E_{r'} = E_t \quad (\text{D.3})$$

and

$$\begin{aligned} H_{II} &:= \sqrt{\frac{\epsilon_0}{\mu_0}} (E_{i'} - E_{r'}) n_1 \cos \theta_1 \\ &= \sqrt{\frac{\epsilon_0}{\mu_0}} E_t n_s \cos \theta_t. \end{aligned} \quad (\text{D.4})$$

Let

$$k_0 h_1 := k_0 (n_1 \Delta_1 \cos \theta_1) \quad (\text{D.5})$$

be the phase difference of the wave between the top and the bottom of the film, with k_0 being the wave number in free space. We can hence rewrite the incidence and reflected waves as:

$$E_{i'} = E_{t'} e^{-ik_0 h_1}$$

and

$$E_{r'} = E_{r''} e^{ik_0 h_1},$$

respectively.

For the two equations at the second interface [(D.3) and (D.4)], we can write:

$$E_{II} = E_{t'} e^{-ik_0 h_1} + E_{r''} e^{ik_0 h_1}$$

and

$$H_{II} = (E_{t'} e^{-ik_0 h_1} - E_{r''} e^{ik_0 h_1}) \sqrt{\frac{\epsilon_0}{\mu_0}} n_1 \cos \theta_1.$$

Solving for $E_{t'}$ and $E_{r''}$, and substituting in (D.1) and (D.2), relates the two interfaces as follows:

$$\begin{aligned} E_I &= E_{II} \cos k_0 h_1 + H_{II} \frac{i \sin k_0 h_1}{\gamma_1} \\ H_I &= E_{II} \gamma_1 i \sin k_0 h_1 + H_{II} \cos k_0 h_1, \end{aligned}$$

where we define

$$\gamma_1 := \sqrt{\frac{\epsilon_0}{\mu_0}} n_1 \cos \theta_1.$$

If we define a matrix \mathbf{M}_1 with

$$\mathbf{M}_1 := \begin{bmatrix} \cos k_0 h_1 & \frac{i \sin k_0 h_1}{\gamma_1} \\ \gamma_1 \sin k_0 h_1 & \cos k_0 h_1 \end{bmatrix},$$

we can write:

$$\begin{pmatrix} E_I \\ H_I \end{pmatrix} = \mathbf{M}_1 \begin{pmatrix} E_{II} \\ H_{II} \end{pmatrix}. \quad (\text{D.6})$$

\mathbf{M}_1 is often referred to as *characteristic matrix* or *transfer matrix*. In order to evaluate the transmissivity and reflectivity of a stack consisting of k film layer (and thus $k + 1$ interfaces), a composition of the corresponding matrices is employed:

$$\begin{pmatrix} E_I \\ H_I \end{pmatrix} = \mathbf{M}_1 \mathbf{M}_2 \dots \mathbf{M}_k \begin{pmatrix} E_{k+1} \\ H_{k+1} \end{pmatrix}.$$

In that fashion, the entire stack or any sub-stack can be evaluated in terms of transmission or reflection behavior.

The previous calculations were conducted for the TE case. The converse TM contribution can be evaluated using the same formulae. The only difference lies in the definition of the γ_j coefficients. While in the TE case, they are defined as

$$\gamma_j^s := \sqrt{\frac{\epsilon_0}{\mu_0}} n_j \cos \theta_j, \quad (\text{D.7})$$

the TM calculation uses

$$\gamma_j^p := \sqrt{\frac{\epsilon_0}{\mu_0}} \frac{n_j}{\cos \theta_j}.$$

In addition, absorption can also be considered, for example, by employing complex refractive indices instead of the real ones demonstrated here.

The transmission and reflection amplitude coefficients can now be evaluated in the following way (here shown for the TE case): The electric and magnetic field components in (D.6) are expressed in terms of the boundary conditions [(D.1), (D.2), (D.4)], giving:

$$\begin{bmatrix} E_i + E_r \\ (E_i - E_r)\gamma_0^s \end{bmatrix} = \mathbf{M} \begin{pmatrix} E_t \\ E_t \gamma_s^s \end{pmatrix}, \quad (\text{D.8})$$

where γ_0^s and γ_s^s are the incidence coefficients for the TE components in the medium and the substrate, respectively, as defined in (D.7). Let

$$r = \frac{E_r}{E_i}$$

be the reflection amplitude coefficient and

$$t = \frac{E_t}{E_i}$$

be the transmission amplitude coefficient, solving (D.8) for r and t yields

$$r = \frac{\gamma_0 m_{11} + \gamma_0 \gamma_s m_{12} - m_{21} - \gamma_s m_{22}}{\gamma_0 m_{11} + \gamma_0 \gamma_s m_{12} + m_{21} + \gamma_s m_{22}} \quad (\text{D.9})$$

$$t = \frac{2\gamma_0}{\gamma_0 m_{11} + \gamma_0 \gamma_s m_{12} + m_{21} + \gamma_s m_{22}} \quad (\text{D.10})$$

$$(\text{D.11})$$

E Parameter Control in Evolutionary Algorithms

This appendix provides a brief overview on the state-of-the-art of parameter control techniques in evolutionary algorithms, complementing the discussion in Section 4.8.¹ As motivated there, parameter control can be broadly categorized as follows (cf. Eiben et al., 1999): (1) *deterministic/dynamic*, (2) *adaptive* and (3) *self-adaptive*.

E.1 Deterministic/Dynamic Techniques

Fogarty (1989) proposed the use of time-varying schedule with an exponential decrease of the mutation rate over time. He also introduced a gene-level mutation scheme, where each bit is tagged with its own mutation rate. Other mutation rate schedules were also reported by Hesser and Männer (1991), Mühlenbein (1992), Bäck (1993) and Yanagiya (1993).

E.2 Adaptive Techniques

Already early versions of evolution strategies (see Section 4.1.1) incorporated an adaptation of mutation rates (Rechenberg, 1973), where the mutation step length was modified according to success (improvement of fitness) of previously performed mutation operations. Little later, self-adaptation, that is, subjecting the step length parameters themselves to selection and mutation, was proposed and proved successful (Schwefel, 1977).

A comprehensive framework of adaptive control parameters in genetic algorithms was devised by Srinivas and Patnaik (1994). Their approach assigns individual crossover and mutation rates to each solution. Above-average solutions receive low rates and are hence “protected;” poorly performing solutions are heavily disrupted by assigning high operator probabilities. Manderick et al. (1991), and similarly Grefenstette (1995), proposed a scheme to select recombination operators. Their approach was to evaluate the impact of the recombination operator on the fitness and choose that operator that was observed best. Another operator selection scheme was developed by Kahng (1995), which was aimed at diversifying the use of different recombination operators. Eshelman and Schaffer (1995) employ a restart scheme in which, if the previous run did not yield any improvement, the restarted GA adopts a different recombination operator. A reward mechanism for operators that are associated with the production of well-performing candidate solutions was proposed by Davis (1989). Later, Julstrom (1995) developed a similar technique with an emphasis on the selection of either mutation or crossover, which in his approach are regarded as competing operators. Wilson (1986) used a scheme in which the entropy of the population served as a measure for the adaptation of crossover probabilities. The main goal was to increase the diversity of solutions. Lee and Takagi (1993) developed meta-evolutionary system based on fuzzy logic. The goal was to generate rules that adapt the control parameters and the population size to a given situation. The input parameters included the relative best and worst fitness values of a population and the fitness change after a control action. The rule set was evolved by applying a GA to a number of test problems.² Afterward, the obtained adaptive GA was tested on a number of

¹Also see the surveys provided by Eiben et al. (1999) and Eiben et al. (2007), and the introduction into self-adaptation by Bäck et al. (2000b).

²The underlying search algorithm was also a genetic algorithm.

benchmark problems, where it was shown to outperform a canonical GA. Another adaptation technique for mutation, crossover and the population size, based on subpopulations, was proposed by Lis and Lis (1996). There, after predefined time intervals, the best-performing GA population determines future modifications of all subpopulations. A similar technique for the adaptation of population sizes was introduced Hinterding et al. (1996).

E.3 Self-adaptive Techniques

Heavily inspired by the techniques used in evolution strategies, Bäck (1992) was one of the first to propose a self-adaption approach for genetic algorithms. His method was later extended and refined by Smith and Fogarty (1996). A similar technique, though using an additional chromosome, was employed by Hinterding (1997) for the selection of one of different mutation operators.

An early self-adaptive approach for crossover was proposed by Schaffer and Morishima (1987). There, the number and location crossover points was encoded and subjected to adaptation. An automaton-based approach, assigning prevalence factors to each allele and so determining the recombination behavior of an operator closely resembling uniform crossover, was shown by White and Oppacher (1994). Another self-adaptive crossover selection scheme was developed by Spears (1995). He added a crossover operator selection parameter, which was attached to the chromosomes and hence simultaneously evolved.

Arabas et al. (1994) propose an implicit adaptive population sizing based on a lifespan of individuals, which is assigned after evaluation and which reflects the proportion of the fitness of the solution and the fitness distribution in the population. Individuals exceeding their lifespan become extinct. This procedure hence yields a dynamically fluctuating population size. The same mechanism was later applied to a steady-state GA by Bäck et al. (2000a). The adaptation of the population size proposed by Smith and Smuda (1995) is aimed at minimizing the selection loss due to sampling error (Goldberg et al., 1992b). A new population size is determined by estimating the overall schema (hyperplane) loss and relating it to a predefined (e.g., user-provided), acceptable schema loss. The next population is increased or decreased depending on the resulting ratio.

Harik and Lobo (1999) present a “parameterless genetic algorithm,” mainly focusing on the adaptation of the population size. They argue that the crossover probability can be eliminated in a straightforward manner, when regarding it as fully correlated with the selection rate. Given both the selection rate s (number of copies of the best solution in the population) and the crossover probability p_c , the growth ratio can be simplified to

$$\gamma = s(1 - p_c).$$

Setting γ to two, ensures that the ratio of the best-observed individuals in the population will grow. Thus by fixing the growth rate, the selection pressure and the crossover rate are determined implicitly. Following Goldberg et al. (1992b) and Harik et al. (1999), the authors devised a routine employing hierarchical subpopulations of incremental size and race between these populations, where smaller sized populations are run in shorter intervals, yet for a shorter time than their larger counterparts. If a small subpopulation has converged or the fitness of a larger population is higher, it is eliminated and recreated (with seeds from the larger population). The main rationales for this approach can be summarized as follows: (1) If a population has converged, the number of samples was too small for an adequate hyperplane processing. (2) Smaller populations exhibiting a lower fitness than their larger counterparts cannot be expected to perform better than the larger population. (3) If a population exhibits a satisfactory maximum fitness, the algorithm can be terminated; no significant computation time is “wasted.” The adaptation of mutation rates was not considered in this concept.

Schlierkamp-Voosen and Mühlenbein (1996) used a meta-level approach by letting subpopulations compete against each other and adapting the population size of each subpopulation according to the competition. For further discussions on parameterless genetic algorithms see Lobo and Lima (2006); Lobo et al. (2007); Bäck et al. (2000a).

Recently, Oltean (2005) applied a (linear) genetic programming approach to evolve optimal genetic algorithms, principally allowing for virtually all control parameters including the flow to be optimized. Promising results on a number of test problems have been reported.

F Benchmark Functions

This appendix provides the formal definitions of the benchmark functions used in this work. More comprehensive discussions on test suites¹ for global optimization approaches are for example given by De Jong (1975, pp. 196–210), Schwefel (1977, pp. 319–354), Törn and Žilinskas (1989, pp. 175–216), Mühlenbein et al. (1991), Whitley et al. (1996) and Digalakis and Margaritis (2001).

F.1 Rosenbrock’s Function

Rosenbrock’s function (Rosenbrock, 1960) features a moderate degree of epistasis and no multimodality. For the purpose of this work, a generalized, multidimensional version of the function is used, which is to be minimized and which is defined as follows:

$$f_2 : \mathbf{x} \mapsto \sum_{i=1}^{N-1} [100(x_{i+1} - x_i^2)^2 + (1 - x_i)^2].$$

The problem can be scaled in a straightforward manner by modifying the dimensionality ($N : 15$ and $N : 30$, in this work). For the (binary-coded) genetic algorithm, the bound-constraints provide an additional factor to alter the complexity ($x_i \in [-10, 10]$ and $x_i \in [-30, 30]$ are used in this work, for example).

In addition to the high dimensionality, the difficulty of Rosenbrock’s function is given by the flat valley around the minimizer $\mathbf{x}^* = \mathbf{1}$. As shown in Figure F.1 for two dimensions, this valley can be pictured as a banana-shaped trench, within which the descent direction to the minimizer is given by small gradients compared to the rest of the search space.

F.2 Shekel’s Foxholes Function

Shekel’s foxholes (Shekel, 1971) functions provide a framework for multimodal problems, with steep local optima and zero gradients on the rest of the search space. In this work, we have employed a less regular maximization variant of the original function in order to reduce the separability of variables—that is, to introduce epistasis. The definition of the function is as follows:

$$f_5 : \mathbf{x} \mapsto \sum_{j=1}^{25} \frac{1}{\frac{1}{j} + |\mathbf{x} - (a)_j|^6},$$

¹From which also the numerical notations for the functions (f_2 : Rosenbrock’s, f_5 : Shekel’s foxholes, f_6 : Rastrigin’s, f_7 : Schwefel’s, f_8 : Griewank’s and f_9 : Ackleys) used here are adopted. Note that f_6 , f_7 are sometimes interchanged.




Figure F.1: Rosenbrock's function in two dimensions on $x_i \in [-2, 2]$, $i = 1, 2$.

where the foxholes $(a)_{ij} = \mathbf{A}$ are defined as:

$$\mathbf{A}^T = \begin{pmatrix} 0 & 0 \\ 1 & 1 \\ 1 & -1.5 \\ -1.5 & 1 \\ -1.1 & -1.3 \\ 2 & 2 \\ -2.1 & 2.7 \\ 2 & -3 \\ -3 & -2.5 \\ 4 & 5 \\ -7 & 4 \\ 3 & -5 \\ -4 & -4 \\ 7 & 6.5 \\ -6.3 & 6.1 \\ 6 & -6 \\ -6 & -6 \\ 0 & 8 \\ -8 & 8 \\ 8 & -8 \\ -8 & 0 \\ 11 & 11 \\ -11 & 11 \\ 11 & -11 \\ -11 & -11 \end{pmatrix}.$$

The global maximizer is hence $\mathbf{x}^* : [-11.0, -11.0]^T$.

The function graph for the domain $[-20, 20]^2$ is depicted in Figure F.2.




Figure F.2: Variant of Shekel's foxholes function on $x_i \in [-20, 20]$, $i = 1, 2$.

F.3 Rastrigin's Function

Rastrigin's function is another highly multimodal function with a multitude of both local maxima and minima, which are both separable with respect to the decision space parameters (cf., e.g., Tang et al., 2010). In other words, Rastrigin's problem does not exhibit any epistasis. The dimensionality of the problem can be arbitrarily scaled (Mühlenbein et al., 1991), yielding the following equation:

$$f_6 : \mathbf{x} \mapsto AN \sum_{i=1}^N [x_i^2 - A \cos(2\pi x_i)],$$

where A is a custom scaling factor, typically set to ten. The search space is normally restricted to the following boundaries: $x_i \in [-5.12, 5.12](i = 1, \dots, N)$.

Figure F.3 shows the function graph in two dimensions. For the benchmark tests in the scope of this work, the 20-dimensional version of the function is to be maximized (typically, a minimization is conducted), where the maximizer is given by $x_i^* \approx -4.523(i = 1, \dots, 20)$.

F.4 Schwefel's Function

Schwefel's multimodal function is defined as (for a discussion, see Mühlenbein et al., 1991):

$$f_7 : \mathbf{x} \mapsto VN - \sum_{i=1}^N x_i \sin(\sqrt{|x_i|}),$$

where $x_i \in [-500, 500](i = 1, \dots, N)$. The parameter V is set to 418.9829 to obtain a function value of about zero for the minimizer $x_i^* : 420.9687(i = 1, \dots, N)$. Where not otherwise noted, a 20-dimensional Schwefel function was used in this work. As shown in Figure F.4, the main difficulty posed by this function is the wide distance of the global minimizer and the secondary optima.




Figure F.3: Rastrigin's function on $x_i \in [-5.12, 5.12]$, $i = 1, 2$.




Figure F.4: Schwefel's function on $x_i \in [-500, 500]$, $i = 1, 2$.

F.5 Griewank's Function

In addition to a high degree of multimodality, Griewank's function (Griewank, 1981) possesses a moderate degree of epistasis. It is defined as:

$$f_8 : \mathbf{x} \mapsto 1 + \frac{1}{4000}|\mathbf{x}|^2 - \prod_{i=1}^N \cos\left(\frac{x_i}{\sqrt{i}}\right),$$

where the search space is restricted to $x_i \in [-660, 600](i = 1, \dots, N)$ and for which the global minimizer is $\mathbf{x}^* : \mathbf{0}$. The 20-dimensional variant was considered in this work. Figure F.5 demonstrates the graph for two dimensions.




Figure F.5: Griewank's function on $x_i \in [-50, 50], i = 1, 2$.

F.6 Ackley's Function

Another common benchmark function was first proposed by Ackley (1987b). It is highly multimodal and separable. Its multidimensional generalization (Bäck, 1996, p. 143) is given by

$$f_9 : \mathbf{x} \mapsto a + e - \exp\left(\frac{1}{N} \sum_{i=1}^N \cos(\omega x_i)\right) - a \exp\left(\frac{-b}{N}|\mathbf{x}|\right),$$

where $a : 20, b : 1/5$ and $\omega : 2\pi$. For the benchmark tests in this work, the minimization—where the global minimizer is $\mathbf{x}^* : \mathbf{0}$ —of the 20-dimensional Ackley function was employed. The search space was restricted to $x_i \in [-30, 30](i = 1, \dots, 20)$.




Figure F.8: Ackley's function in two-dimensions; (a) Search space as used for the benchmark tests $x_i \in [-30, 30]$, $i = 1, 2$. (b) Close-up revealing the multimodal structure of the function.

Bibliography

- E. Abbe. Beiträge zur Theorie des Mikroskops und der mikroskopischen Wahrnehmung. *Archiv für mikroskopische Anatomie*, 9(1):413–418, 1873.
- D. S. Abrams and L. Pang. Fast inverse lithography technology. In D. G. Flagello (editor), *Proc. SPIE*, volume 6154, page 61541J. SPIE, 2006.
- D. Ackley. Stochastic iterated genetic hillclimbing. 1987a.
- D. H. Ackley. *A connectionist machine for genetic hillclimbing*, volume 28. Kluwer Academic Publishers, Boston, MA, 1987b.
- K. Adam. *Domain Decomposition Methods for the Electromagnetic Simulation of Scattering from Three-Dimensional Structures with Applications in Lithography*. Ph.D. dissertation, University of California, Berkeley, 2001.
- K. Adam, M. C. Lam, N. Cobb and O. Touban. Application of the hybrid Hopkins–Abbe method in full-chip OPC. *Microelectronic Engineering*, 86(4-6):492–496, 2009.
- S. Adams, J. Payne and R. Boppana. Finite Difference Time Domain (FDTD) Simulations Using Graphics Processors. *High Performance Computing Modernization Program Users Group Conference*, pages 334–338, 2007.
- C.-N. Ahn, K.-H. Baik, Y.-S. Lee, H.-E. Kim, I. Hur, Y. Kim, J.-H. Kim and S.-H. Choi. Study of optical proximity effects using off-axis illumination with attenuated phase shift mask. In T. A. Brunner (editor), *Proceedings of SPIE*, volume 2440, pages 222–239. SPIE, 1995.
- C.-N. Ahn, H.-B. Kim and B. Ki-Ho. Novel approximate model for resist process. In *Proceedings of SPIE*, pages 752–763. SPIE, 1998.
- M. M. Alkaisi, R. J. Blaikie, S. J. McNab, R. Cheung and D. R. S. Cumming. Sub-diffraction-limited patterning using evanescent near-field optical lithography. *Applied Physics Letters*, 75(22):3560, 1999.
- P. Allen. Laser scanning for semiconductor mask pattern generation. *Proceedings of the IEEE*, 90(10):1653–1669, 2002.
- R. Allenson. Genetic algorithms with gender for multi-function optimisation. *Edinburgh Parallel Comput. Center, Edinburgh, Scotland, Tech. Rep*, pages 92–01, 1992.
- L. Altenberg. The schema theorem and Price’s theorem. *Foundations of genetic algorithms*, 3:23–49, 1995.
- J. Andersson and D. Wallace. Pareto optimization using the struggle genetic crowding algorithm. *Engineering Optimization*, 34(6):623–643, 2002.
- S. Ando, J. Sakuma and S. Kobayashi. Adaptive isolation model using data clustering for multimodal function optimization. In *Proceedings of the 2005 Conference on Genetic and Evolutionary Computation*, pages 1417–1424. ACM Press, New York, New York, USA, 2005.
- T. Ando, M. Takeshita, R. Takasu, Y. Yoshii, J. Iwashita, S. Matsumaru, S. Abe and T. Iwai. Pattern freezing process free litho-litho-etch double patterning. In *Proceedings of SPIE*, volume 7140, page 71402H. 2008.

- D. Andre, F. H. Bennett III and J. R. Koza. Discovery by Genetic Programming of a Cellular Automata Rule that is Better than any Known Rule for the Majority Classification Problem. In J. R. Koza, D. E. Goldberg, D. B. Fogel and R. L. Riolo (editors), *Genetic Programming 1996: Proceedings of the First Annual Conference*, pages 3–11. MIT Press, Stanford University, CA, USA, 1996.
- P. J. Angeline. *Two self-adaptive crossover operators for genetic programming*, pages 89–109. The MIT Press, Cambridge, MA, USA, 1996.
- P. J. Angeline. Subtree Crossover: Building Block Engine or Macromutation? In J. R. Koza, K. Deb, M. Dorigo, D. B. Fogel, M. Garzon, H. Iba and R. L. Riolo (editors), *Genetic Programming 1997: Proceedings of the Second Annual Conference*, pages 9–17. Morgan Kaufmann, Stanford University, CA, USA, 1997.
- P. J. Angeline and J. B. Pollack. The evolutionary induction of subroutines. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*, pages 236–241. Lawrence Erlbaum, 1992.
- J. Antonisse. A New Interpretation of Schema Notation that Overtums the Binary Encoding Constraint. pages 86–91. Morgan Kaufmann Publishers Inc., 1989.
- J. Arabas, Z. Michalewicz and J. Mulawka. GAVaPS-a genetic algorithm with varying population size. In *Evolutionary Computation, 1994. IEEE World Congress on Computational Intelligence., Proceedings of the First IEEE Conference on*, pages 73–78. IEEE, 1994.
- S. Areibi and Z. Yang. Effective Memetic Algorithms for VLSI design = Genetic Algorithms + local search + multi-level clustering. *Evolutionary computation*, 12(3):327–53, 2004.
- M. Asano. Sub-100 nm lithography with KrF exposure using multiple development method. *Jpn. J. Appl. Phys*, 38(12B pt 1):6999–6999, 1999.
- T. Bäck. Self-adaptation in genetic algorithms. In *Towards a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*, pages 263–271. 1992.
- T. Bäck. Optimal mutation rates in genetic search. In *Proceedings of the Fifth International Conference on Genetic Algorithms*, pages 2–8. 1993.
- T. Bäck. Selective pressure in evolutionary algorithms: a characterization of selection mechanisms. In *Proceedings of the First IEEE Conference on Evolutionary Computation. IEEE World Congress on Computational Intelligence*, pages 57–62. IEEE, 1994.
- T. Bäck. Generalized Convergence Models for Tournament- and (μ , λ)-Selection. In *Proceedings of the 6th International Conference on Genetic Algorithms*, pages 2–8. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1995.
- T. Bäck. *Evolutionary Algorithms in Theory and Practice*, volume 2. Oxford University Press, 1996.
- T. Bäck. Self-adaptation. *Evolutionary Computation* 2, 2:188–211, 2000.
- T. Bäck, A. Eiben and N. A. L. van der Vaart. An Empirical Study on GAs "Without Parameters". In M. Schoenauer, K. Deb, G. Rudolph, X. Yao, E. Lutton, J. J. M. Guervós and H.-P. Schwefel (editors), *PPSN*, volume 1917 of *Lecture Notes in Computer Science*, pages 315–324. Springer, 2000a.
- T. Bäck, D. Fogel and Z. Michalewicz (editors). *Evolutionary Computation 2: Advanced Algorithms and Operators*. Number v. 1 in *Evolutionary computation*. Institute of Physics Publishing, 2000b.
- T. Bäck, G. Rudolph and H.-P. Schwefel. Evolutionary programming and evolution strategies: Similarities and differences. In *Proceedings of the Second Annual Conference on Evolutionary Programming*, pages 11–22. 1993.

- J. Bader and E. Zitzler. HypE: an algorithm for fast hypervolume-based many-objective optimization. *Evolutionary computation*, 19(1):45–76, 2011.
- J. D. Bagley. *The behavior of adaptive systems which employ genetic and correlation algorithms*. Ph.D. dissertation, University of Michigan, Ann Arbor, MI, USA, Ann Arbor, MI, USA, 1967.
- S. Bajt, H. Chapman, N. Nguyen, J. Alameda, J. Robinson, M. Malinowski, E. Gullikson, A. Aquila, C. Tarrio and S. Grantham. Design and performance of capping layers for EUV multilayer mirrors. In *Proceedings of SPIE*, volume 5037, page 236. 2003.
- S. Bajt and D. G. Stearns. High temperature stability multilayers for EUV condenser optics. *Applied Optics*, 44(UCRL-JRNL-212020), 2005.
- J. Baker. Adaptive selection methods for genetic algorithms. In *Proceedings of the 1st International Conference on Genetic Algorithms*, pages 101–111. L. Erlbaum Associates Inc., 1985.
- J. E. Baker. Reducing bias and inefficiency in the selection algorithm. In *Proceedings of the Second International Conference on Genetic Algorithms on Genetic algorithms and their application*, pages 14–21. L. Erlbaum Associates Inc., 1987.
- V. Bakshi (editor). *EUV Sources for Lithography*. SPIE, 1000 20th Street, Bellingham, WA 98227-0010 USA, 2006.
- J. M. Baldwin. A New Factor in Evolution. *The American Naturalist*, 30(354):441–451, 1896.
- V. Banine, J. Benschop, M. Leenders and R. Moors. Relationship between an EUV source and the performance of an EUV lithographic system. In *Proceedings of SPIE*, volume 3997, page 126. SPIE, 2000.
- W. Banzhaf. The “molecular” traveling salesman. *Biological Cybernetics*, 64(1):7–14, 1990.
- W. Banzhaf. Genetic programming for pedestrians. In *Proceedings of the 5th International Conference on Genetic Algorithms, ICGA-93*, page 628. 1993.
- G. Bao. Finite Element Approximation of Time Harmonic Waves in Periodic Structures. *SIAM Journal on Numerical Analysis*, 32(4):1155–1169, 1995.
- A. Barenbaum. *Optimierung von Polarisationsplittern*. Diplomarbeit, University of Erlangen/Nuremberg, 2010.
- J. Bartelt. Masked Ion Beam Lithography: An Emerging Technology. *Solid State Technol*, 29:215, 1986.
- A. Barty, S. Hau-Riege, D. G. Stearns, M. Clift, P. Mirkarimi, E. Gullikson, H. Chapman and D. Sweeney. A method for repairing amplitude defects in multilayer-coated EUV mask blanks. *Applied Optics*, 43(UCRL-JRNL-200816), 2003.
- A. Barty, Y. Liu, E. Gullikson, J. Taylor and O. Wood. Actinic inspection of multilayer defects on EUV masks. In *Proceedings of SPIE*, volume 5751, page 651. SPIE, 2005.
- T. Battiti. First- and Second-Order Methods for Learning: Between Steepest Descent and Newton’s Method. *Neural Computation*, 4(2):141–166, 1992.
- D. Battle and M. D. Vose. Isomorphisms of genetic algorithms. *Artificial Intelligence*, 60(1):155–165, 1993.
- D. Beasley, D. Bull and R. Martin. A sequential niche technique for multimodal function optimization. *Evolutionary computation*, 1(2):101–125, 1993.

- H. W. Becker, J. Chey, F. Sobel, F. Schmidt, M. Renno, U. Buttgereit, M. Angelopoulos, G. Hess and Konrad Knapp. Development of a new PSM film system for 157-nm extensible to high-transmission 193 nm lithography. In *Proceedings of SPIE*, volume 5256, pages 204–212. SPIE, 2003.
- M. J. Beesley and J. G. Castledine. The Use of Photoresist as a Holographic Recording Medium. *Applied Optics*, 9(12):2720, 1970.
- N. Bekiaris, H. Cervera, J. Dai, R.-h. Kim, A. Acheta, T. Wallow, J. Kye, H. J. Levinson, T. Nowak and J. Yu. A lithographic and process assessment of photoresist stabilization for double-patterning using 172-nm photoresist curing. In *Proceedings of SPIE*, volume 6923, page 692321. SPIE, 2008.
- R. Bellman. *Adaptive Control Processes: A Guided Tour*, volume 8. Princeton University Press, 1961.
- C. Bencher, Y. Chen, H. Dai, W. Montgomery and L. Huli. 22nm half-pitch patterning by CVD spacer self alignment double patterning (SADP). In *Proceedings of SPIE*, volume 6924, pages 69244E–69244E–7. SPIE, 2008.
- G. Beni and J. Wang. Swarm intelligence in cellular robotic systems. In *Proceedings of the NATO Advanced Workshop on Robots and Biological Systems*. 1989.
- N. Benoit, S. Yulin, T. Feigl and N. Kaiser. Radiation stability of EUV Mo/Si multilayer mirrors. *Physica B: Condensed Matter*, 357(1-2):222–226, 2005.
- J.-P. Berenger. A perfectly matched layer for the absorption of electromagnetic waves. *Journal of Computational Physics*, 114(2):185–200, 1994.
- R. Berger, J. Kauffmann, N. Kerwien, W. Osten and H. Tiziani. Rigorose Beugungssimulation: Ein Vergleich zwischen RCWA, FDTD und der Finiten Elemente Methode. *DGaO-Proceedings*, 2004.
- K. Berggren, A. Bard, J. Wilbur, J. Gillaspay, A. Helg, J. McClelland, S. Rolston, W. Phillips, M. Prentiss and G. Whitesides. Microlithography by using neutral metastable atoms and self-assembled monolayers. *Science*, 269(5228):1255–1257, 1995.
- L. Berke and P. Hajela. Applications of artificial neural nets in structural mechanics. *Structural Optimization*, 4(2):90–98, 1992.
- D. Bernard. Simulation of focus effects in photolithography. *IEEE Transactions on Semiconductor Manufacturing*, 1(3):85–97, 1988.
- A. Berry and P. Vamplew. PoD Can Mutate: A simple dynamic directed mutation approach for genetic algorithms. In *AISAT2004: International Conference on Artificial Intelligence in Science and Technology*. 2004.
- H. Bersini and F. Varela. Hints for adaptive problem solving gleaned from immune networks. In *Parallel problem solving from nature*, pages 343–354. Springer, 1991.
- M. Besacier and P. Schiavone. Shadowing effect minimization in EUV mask by modeling. In *Proceedings of SPIE*, volume 5446, page 849. 2004.
- A. D. Bethke. *Genetic algorithms as function optimizers*. Ph.D. dissertation, University of Michigan, 1981.
- E. Betzig and J. K. Trautman. Near-field optics: microscopy, spectroscopy, and surface modification beyond the diffraction limit. *Science*, 257(5067):189–95, 1992.
- N. Beume, B. Naujoks and M. Emmerich. SMS-EMOA: Multiobjective selection based on dominated hypervolume. *European Journal of Operational Research*, 181(3):1653–1669, 2007.

- J. J. Biafore, M. D. Smith, T. Wallow, P. Nalleau, D. Blankenship and Y. Deng. Pattern prediction in EUV resists. *Library*, 7520:75201P–75201P–13, 2009.
- P. Björnängen, M. Ekberg, T. Öström, H. Fosshaug, J. Karlsson, C. Björnberg, F. Nikolaieff and M. Karlsson. DOE manufacture with the DUV SLM-based Sigma7300 laser pattern generator. In B. W. Smith (editor), *Proceedings of SPIE*, volume 5377, pages 1866–1875. SPIE, 2004.
- R. Blaikie, D. O. S. Melville and M. Alkai. Super-resolution near-field lithography using planar silver lenses: A review of recent developments. *Microelectronic Engineering*, 83(4-9):723–729, 2006.
- J. Blaschak and G. Kriegsmann. A comparative study of absorbing boundary conditions. *Journal of Computational Physics*, 77(1):109–139, 1988.
- T. Blickle and L. Thiele. A mathematical analysis of tournament selection. In *Proceedings of the Sixth International Conference on Genetic Algorithms*, pages 9–16. 1995.
- T. Blickle and L. Thiele. A comparison of selection schemes used in genetic algorithms. *Evolutionary Computation*, 4(4):311–347, 1997.
- P. T. Boggs and J. W. Tolle. Sequential Quadratic Programming. *Acta Numerica*, 4:1–51, 1995.
- B. Bollepalli, M. Khan and F. Cerrina. Imaging properties of the extreme ultraviolet mask. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 16:3444, 1998.
- L. Booker. Improving search in genetic algorithms. *Genetic algorithms and simulated annealing*, 61:73, 1987.
- M. Born and E. Wolf. *Principles of optics: electromagnetic theory of propagation, interference and diffraction of light*. Cambridge University Press, 7 edition, 1999.
- Y. Borodovsky, W.-H. Cheng, R. Schenker and V. Singh. Pixelated phase mask as novel lithography RET. In *Proceedings of SPIE*, volume 6924, pages 69240E–69240E–14. 2008.
- J. W. Bossung. Projection Printing Characterization. In J. W. Giffin (editor), *Proceedings of SPIE*, volume 100 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pages 80–85. 1977.
- M. J. Box. A New Method of Constrained Optimization and a Comparison With Other Methods. *The Computer Journal*, 8(1):42–52, 1965.
- R. M. Brady. Optimization strategies gleaned from biological evolution. *Nature*, 317(6040):804–806, 1985.
- M. F. Bramlette. Initialization, Mutation and Selection Methods in Genetic Algorithms for Function Optimization. In *Proceedings of the Fourth International Conference on Genetic Algorithms*, page 100. Morgan Kaufmann Publishers, 1991.
- D. Brandt, I. Fomenkov, A. Ershov, W. Partlo, D. Myers, N. Böwering, N. Farrar, G. Vaschenko, O. Khodykin, A. Bykanov and Others. LPP source system development for HVM. In *Proceedings of SPIE*, volume 7271, page 727103. 2009.
- J. Branke, H. Schmeck and Others. Guidance in evolutionary multi-objective optimization. *Advances in Engineering Software*, 32(6):499–507, 2001.
- S. Braun, H. Mai, M. Moss, R. Scholz and A. Leson. Mo/Si Multilayers with Different Barrier Layers for Applications as Extreme Ultraviolet Mirrors. *Japanese Journal of Applied Physics*, 41(Part 1, No. 6B):4074–4081, 2002.

- C. L. Bridges and D. E. Goldberg. An analysis of reproduction and crossover in a binary-coded genetic algorithm. In *Proceedings of the Second International Conference on Genetic Algorithms on Genetic algorithms and their application*, pages 9–13. L. Erlbaum Associates Inc., Hillsdale, NJ, USA, 1987.
- A. Brindle. *Genetic algorithms for function optimization*. Ph.D. dissertation, University of Alberta, 1981.
- T. Brist and J. Torres. Model-assisted placement of subresolution assist features: experimental results. In *Proceedings of SPIE*, volume 5042, page 99. SPIE, 2003.
- D. Brockhoff, T. Friedrich and F. Neumann. Analyzing hypervolume indicator based algorithms. In *Parallel Problem Solving from Nature (PPSN X)*, pages 651–660. Springer, 2008.
- D. J. V. D. Broeke, J. F. Chen, T. L. Laidig, S. Hsu, K. E. Wampler, R. J. Socha and J. S. Petersen. Complex 2D pattern lithography at $\lambda/4$ resolution using chromeless phase lithography (CPL). In *Proceedings of SPIE*, volume 4691, pages 196–214. SPIE, 2002.
- C. Broyden. Quasi-Newton methods and their application to function minimization. *Mathematics of Computation*, 21:368–381, 1967.
- T. A. Brunner. Rim phase-shift mask combined with off-axis illumination: a path to $.5\lambda/\text{NA}$ geometries. In J. D. Cuthbert (editor), *Proceedings of SPIE*, volume 1927, pages 54–62. 1993.
- T. A. Brunner. Impact of lens aberrations on optical lithography. *IBM Journal of Research and Development*, 41(1):57–67, 1997.
- T. A. Brunner, N. Seong, W. D. Hinsberg, J. A. Hoffnagle, F. A. Houle and M. I. Sanchez. High numerical aperture lithographic imagery at the Brewster angle. *Journal of Microlithography, Microfabrication, and Microsystems*, 1(3):188–196, 2002.
- K. Bubke, M. Sczyrba and C. Pierrat. Applicability of alternating phase shifting masks using polarized light. In *Proceedings of SPIE*, volume 6154, pages 615449–615449–11. SPIE, 2006.
- L. Bull. On Model-Based Evolutionary Computation. *Soft Computing - A Fusion of Foundations, Methodologies and Applications*, 3(2):76–82, 1999.
- S. Burger, L. Zschiedrich, F. Schmidt, P. Evanschitzky and A. Erdmann. Benchmark of rigorous methods for electromagnetic field simulations. In H. Kawahira and L. S. Zurbrick (editors), *Proceedings of SPIE*, volume 7122, pages 71221S–71221S–12. 2008.
- M. Burkhardt, A. Yen, C. Progler and G. Wells. Illuminator design for the printing of regular contact patterns. *Microelectronic Engineering*, 41-42:91–95, 1998.
- J. H. Burnett, S. G. Kaplan, E. L. Shirley, P. J. Tompkins and J. E. Webb. High-index materials for 193 nm immersion lithography. In *Proceedings of SPIE*, volume 5754, pages 611–621. 2005.
- S. D. Burns, A. Gardiner, V. Krukonis, P. Wetmore, J. Lutkenhaus, G. Schmid, L. Flanagan and C. Willson. Understanding nonlinear dissolution rates in photoresists. In *Proceedings of SPIE*, volume 4345, pages 37–49. SPIE, 2001.
- J. D. Byers, S. Lee, K. Jen, P. A. Zimmerman, N. Turro and C. G. Willson. Double exposure materials: simulation study of feasibility. *Journal of Photopolymer Science and Technology*, 20(5):707–717, 2007.
- J. D. Byers, M. Smith and C. A. Mack. 3D lumped parameter model for lithographic simulations. In *Proceedings of SPIE*, pages 125–137. SPIE, 2002.
- J. D. Byers, M. Smith and C. A. Mack. Lumped parameter model for chemically amplified resists. In *Proceedings of SPIE*, pages 1462–1474. SPIE, 2004.

- J. D. Byers, M. Smith, C. A. Mack and J. Biafore. Modeling soft-bake effects in chemically amplified resists. *Proceedings of SPIE*, 5039:1143–1154, 2003.
- R. H. Byrd and J. Nocedal. An analysis of reduced Hessian methods for constrained optimization. *Mathematical Programming*, 49:285–323, 1990.
- R. H. Byrd and R. Schnabel. A trust region algorithm for nonlinearly constrained optimization. *SIAM Journal on Numerical Analysis*, 24(5):1152–1170, 1987.
- S. Cahon, N. Melab and E.-G. Talbi. ParadisEO: A Framework for the Reusable Design of Parallel and Distributed Metaheuristics. *Journal of Heuristics*, 10(3):357–380, 2004.
- M. Campbell, D. Sharp, M. Harrison, R. Denning and A. Turberfield. Fabrication of photonic crystals for the visible spectrum by holographic lithography. *Nature*, 404(6773):53–6, 2000.
- E. Cantú-Paz. A survey of parallel genetic algorithms. Illigal report no. 97003, University of Illinois at Urbana-Champaign, 1997.
- L. Capodiecici. *Optimization techniques for VLSI process modeling and TCAD in semiconductor manufacturing*. Ph.D. dissertation, University of Wisconsin, 1997.
- L. Capodiecici, A. Krasnoperova, F. Cerrina, C. Lyons, C. Spence and K. Early. Novel postexposure bake simulator: First results. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 13(6):2963–2967, 1995.
- P. Carcia, R. French, K. Sharp, J. Meth and B. Smith. Materials screening for attenuating embedded phase-shift photoblanks for DUV and 193 nm photolithography. In *Sixteenth Annual BACUS Symposium on Photomask Technology and Management*, GV Shelden and JA Reynolds, ed., *Proc. SPIE*, volume 2884, pages 255–263. 1996.
- W. C. Carpenter and J. F. M. Barthelemy. A comparison of polynomial approximations and artificial neural nets as response surfaces. *Structural Optimization*, 5(3):166–174, 1993.
- J. Carriere and T. Hudson. Diffractive Optics Enable Improved Illumination, Resolution. *Semiconductor international*, 32(13):12–15, 2009.
- R. Caruana and J. D. Schaffer. Representation and hidden bias: Gray vs. binary coding for genetic algorithms. In *Proceedings of the Fifth International Conference on Machine Learning*, pages 153–161. Ann Arbor, Michigan, USA, 1988.
- R. A. Caruana, L. J. Eshelman and J. D. Schaffer. Representation and hidden bias II: eliminating defining length bias in genetic search via shuffle crossover. In *Proceedings of the 11th international joint conference on Artificial intelligence - Volume 1*, pages 750–755. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1989.
- W. Cedeño, V. R. Vemuri and T. Slezak. Multiniche crowding in genetic algorithms and its application to the assembly of DNA restriction-fragments. *Evolutionary Computation*, 2(4):321–345, 1994.
- M. R. Celis, J. E. J. Dennis and R. A. Tapia. A trust region strategy for nonlinear equality constrained optimization. *Numerical Optimization*, 1994:71–82, 1985.
- F. Cerrina, S. Bollepalli, M. Khan, H. H. Solak, W. Li and D. He. Image formation in EUV lithography: Multilayer and resist properties. *Microelectronic Engineering*, 53(1-4):13–20, 2000.
- N. Chai-ead, P. Aungkulanon and P. Luangpaiboon. Bees and Firefly Algorithms for Noisy Non-Linear Optimisation Problems. In *Proceedings of the International MultiConference of Engineers and Computer Scientists*, volume 2. 2011.

- U. K. Chakraborty (editor). *Advances in Differential Evolution*, volume 143 of *Studies in Computational Intelligence*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.
- T. H. P. Chang. Proximity effect in electron-beam lithography. *Journal of Vacuum Science and Technology*, 12(6):1271, 1975.
- W. Chang, E. Kim, Y. Kang, S. Park, C. Lim, K. Won, J. Kim and H.-K. Oh. Characteristics and prevention of pattern collapse in EUV lithography. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 6517, page 85. SPIE, 2007.
- A. Charnes and W. Cooper. Optimal estimation of executive compensation by linear programming. *Management science*, 1(2):138–151, 1955.
- A. Charnes and W. W. Cooper. *Management Models and Industrial Applications of Linear Programming: v. 1*. John Wiley & Sons Inc, 1961.
- C. C. P. Chen, A. Gurhanli, T.-Y. Chiang, J.-J. Hong and L. S. Melvin. Abbe singular-value decomposition: Compact Abbe's kernel generation for microlithography aerial image simulation using singular-value decomposition method. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 26(6):2322, 2008.
- F. Chen. Asymmetry and thickness effects in reflective EUV masks. In *Proceedings of SPIE*, volume 5037, page 347. 2003.
- H. L. Chen, H. C. Cheng, T. S. Ko, F. H. Ko and T. C. Chu. High reflectance of reflective-type attenuated-phase-shifting masks for extreme ultraviolet lithography with high inspection contrast in deep ultraviolet regimes. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 22(6):3049, 2004.
- J. Chen, T. Laidig, K. Wampler and R. Caldwell. Optical proximity correction for intermediate-pitch features using sub-resolution scattering bars. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 15:2426, 1997.
- J. F. Chen, J. S. Petersen, R. J. Socha, T. L. Laidig, K. E. Wampler, K. H. Nakagawa, G. P. Hughes, S. S. MacDonald and W. Ng. Binary halftone chromeless PSM technology for $\lambda/4$ optical lithography. In *Proceedings of SPIE*, volume 4346, pages 515–533. SPIE, 2001.
- Y. Chen, K. Wu, Z. Shi and X. Yan. A feasible model-based OPC algorithm using Jacobian matrix of intensity distribution functions. In *Proceedings of SPIE*, volume 6520, pages 65204C–1. 2007.
- W. Cheng, J. Farnsworth, W. Kwok, A. Jamieson, N. Wilcox, M. Vernon, K. Yung, Y. Liu, J. Kim, E. Frendberg and Others. Fabrication of defect-free full-field pixelated phase mask. In *Proceedings of SPIE*, volume 6924, page 69241G. SPIE, 2008.
- A. Chiba, M. Sugawara, H. Yamanashi and I. Nishiyama. Assessment of Heat Deformation and Throughput for Selecting Mask Substrate Material for Extreme Ultraviolet Lithography. *Japanese Journal of Applied Physics*, 41(Part 1, No. 11A):6498–6505, 2002.
- S.-J. Choi, H.-S. Cha, S.-Y. Yoon, S.-M. Jung, S.-S. Choi and S.-H. Jeong. Dry etching characteristics of attenuated phase-shifting masks using Cl. In *Proceedings of SPIE*, volume 4562, pages 561–570. SPIE, 2002.
- C.-W. Chu, B. Tsao, K. Chiou, S. Lee, J. Huang, Y. Liu, T. Lin, A. Moore and L. Pang. Enhancing DRAM printing process window by using inverse lithography technology (ILT). In D. G. Flagello (editor), *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 6154 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pages 1231–1240. 2006.

- C. H. Clifford, T. Chan and A. Neureuther. Compensation methods for buried defects in extreme ultraviolet lithography masks. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 7636, page 66. 2010.
- C. H. Clifford and A. R. Neureuther. Fast simulation methods and modeling for extreme ultraviolet masks with buried defects. *Journal of Micro/Nanolithography, MEMS and MOEMS*, 8(3):031402, 2009.
- N. Cobb. *Fast optical and process proximity correction algorithms for integrated circuit manufacturing*. Ph.D. dissertation, University of California, Berkeley, 1998.
- N. Cobb and Y. Granik. Model-based OPC using the MEEF matrix. In B. J. Grenon and K. R. Kimmel (editors), *Proceedings of SPIE*, volume 4889, pages 1281–1292. SPIE, Monterey, CA, USA, 2002.
- N. Cobb, A. Zakhor and E. Miloslavsky. Mathematical and CAD framework for proximity correction. In *Proceedings of SPIE*, volume 2726, page 208. 1996.
- C. A. Coello Coello. Theoretical and numerical constraint-handling techniques used with evolutionary algorithms: a survey of the state of the art. *Computer methods in applied mechanics and engineering*, 191(11-12):1245–1287, 2002.
- C. A. Coello Coello, G. Lamont and D. A. Van Veldhuizen. *Evolutionary algorithms for solving multi-objective problems*, volume 5 of *Genetic and Evolutionary Computation Series*. Springer-Verlag New York Inc, Boston, MA, 2007.
- J. P. Cohoon, S. U. Hegde, W. N. Martin and D. Richards. Punctuated equilibria: a parallel genetic algorithm. In *Proceedings of the Second International Conference on Genetic Algorithms on Genetic algorithms and their application*, pages 148–154. L. Erlbaum Associates Inc., Hillsdale, NJ, USA, 1987.
- E. Collingwood. *Investigation of a Multiple Chromosome Evolutionary Algorithm for Bus Driver Scheduling and Other Problems*. M.sc. thesis, 1995.
- A. R. Conn, N. I. M. Gould and P. L. Toint. Convergence of quasi-Newton matrices generated by the symmetric rank one update. *Mathematical Programming*, 50(1-3):177–195, 1991.
- D. W. Corne, N. R. Jerram, J. D. Knowles and M. J. Oates. PESA-II: Region-based Selection in Evolutionary Multiobjective Optimization. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO'2001)*, volume 5717, pages 283–290. Morgan Kaufmann, 2001.
- D. W. Corne and J. D. Knowles. No free lunch and free leftovers theorems for multiobjective optimization problems. In *Evolutionary Multi-Criterion Optimization*, pages 66–66. Springer, 2003a.
- D. W. Corne and J. D. Knowles. Some multiobjective optimizers are better than others. In *The 2003 Congress on Evolutionary Computation, 2003. CEC '03.*, volume 4, pages 2506–2512. IEEE, 2003b.
- D. W. Corne, J. D. Knowles and M. Oates. The Pareto envelope-based selection algorithm for multi-objective optimization. In *Parallel Problem Solving from Nature PPSN VI*, pages 839–848. Springer, 2000.
- P. Cowling, G. Kendall and E. Soubeiga. A Hyperheuristic Approach to Scheduling a Sales Summit. *Practice and Theory of Automated Timetabling III*, pages 176–190, 2001.
- J. C. Culberson. On the Futility of Blind Search: An Algorithmic View of “No Free Lunch”. *Evolutionary Computation*, 6(2):109–127, 1998.
- V. Cutello and G. Nicosia. An immunological approach to combinatorial optimization problems. In *Advances in Artificial Intelligence—IBERAMIA 2002*, pages 361–370. Springer, 2002.

- G. Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals, and Systems (MCSS)*, 2(4):303–314, 1989.
- T. Dam, V. Tolani, P. Hu, K.-H. Baik, L. Pang, B. Gleason, S. D. Slonaker and J. K. Tyminski. Source-mask optimization (SMO): from theory to practice. In *Proc. SPIE*, volume 7640, pages 764028–764028–10. 2010.
- C. Darwin. *The origin of species by means of natural selection: or, the preservation of favored races in the struggle for life*. Modern Library (Random House), New York, 1998.
- I. Das. Normal-boundary intersection: A new method for generating the Pareto surface in nonlinear multicriteria optimization problems. *SIAM Journal on Optimization*, 1998.
- I. Das and J. E. Dennis. A closer look at drawbacks of minimizing weighted sums of objectives for Pareto set generation in multicriteria optimization problems. *Structural Optimization*, 14(1):63–69, 1997.
- D. Dasgupta and L. Niño. *Immunological computation: theory and applications*. Auerbach Publications, 2008.
- S. Dasgupta, S. Das, A. Biswas and A. Abraham. On stability and convergence of the population-dynamics in differential evolution. *AI Communications*, 22(1):1–20, 2009.
- Y. Davidor. Epistasis variance: Suitability of a representation to genetic algorithms. *Complex Systems*, 4(4):369–383, 1990.
- Y. Davidor. A naturally occurring niche & species phenomenon: The model and first results. In *Proceedings of the Fourth International Conference on Genetic Algorithms*, volume 257, page 263. Morgan Kaufmann Publishers, 1991.
- P. Davids and S. Bollepalli. Generalized inverse problem for partially coherent projection lithography. In *Proceedings of SPIE*, volume 6924, page 69240X. 2008.
- L. D. Davis. Job shop scheduling with genetic algorithms. In *Proceedings of the 1st International Conference on Genetic Algorithms*, pages 136–140. L. Erlbaum Associates Inc., 1985.
- L. D. Davis. Adapting Operator Probabilities in Genetic Algorithms. In *Proceedings of the 3rd International Conference on Genetic Algorithms*, pages 61–69. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1989.
- L. D. Davis. *Handbook Of Genetic Algorithms*. Van Nostrand Reinhold Company, 1991.
- L. D. Davis and S. Coombs. Genetic algorithms and communication link speed design: theoretical considerations. In *Proceedings of the Second International Conference on Genetic Algorithms on Genetic algorithms and their application*, pages 252–256. L. Erlbaum Associates Inc., 1987.
- R. Dawkins. *The Selfish Gene*. Oxford University Press, 1976.
- P. de Bisschop, A. Erdmann and A. Rathsfeld. Simulation of the effect of a resist-surface bound air bubble on imaging in immersion lithography. In B. W. Smith (editor), *Optical Microlithography XVIII*, volume 5754, pages 243–253. SPIE, 2004.
- L. De Castro and F. Von Zuben. The clonal selection algorithm with engineering applications. In *Proceedings of GECCO*, volume 2000, pages 36–39. 2000.
- K. A. De Jong. *Analysis of the behavior of a class of genetic adaptive systems*. Ph.D. dissertation, University of Michigan, Ann Arbor, MI, 1975.

- K. A. De Jong. Genetic algorithms are NOT function optimizers. *Foundations of Genetic Algorithms*, 2:5–17, 1993.
- K. A. De Jong and W. M. Spears. A formal analysis of the role of multi-point crossover in genetic algorithms. *Annals of Mathematics and Artificial Intelligence*, 5(1):1–26, 1992.
- K. Deb. *Genetic algorithms in multimodal function optimization*. Master thesis, University of Alabama, 1989.
- K. Deb. Multi-objective genetic algorithms: Problem difficulties and construction of test problems. *Evolutionary computation*, 7(3):205–230, 1999.
- K. Deb. *Multi-Objective Optimization Using Evolutionary Algorithms*. John Wiley & Sons, Inc., New York, NY, USA, 2001.
- K. Deb and R. B. Agrawal. Simulated Binary Crossover for Continuous Search Space. Technical report, Departement of Mechanical Engineering, Indian Institute of Technology, Kanpur, India, 1994.
- K. Deb and A. Anand. A computationally efficient evolutionary algorithm for real-parameter optimization. *Evolutionary computation*, 2002.
- K. Deb and D. E. Goldberg. An Investigation of Niche and Species Formation in Genetic Function Optimization. pages 42–50. Morgan Kaufmann Publishers Inc., 1989.
- K. Deb, A. Pratap, S. Agarwal and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197, 2002a.
- K. Deb, A. Sinha and S. Kukkonen. Multi-objective test problems, linkages, and evolutionary methodologies. In *Proceedings of the 8th Annual Conference on Genetic and Evolutionary Computation*, pages 1141–1148. ACM Press, New York, New York, USA, 2006.
- K. Deb and J. Sundar. Reference point based multi-objective optimization using evolutionary algorithms. In *Proceedings of the 8th annual conference on Genetic and evolutionary computation - GECCO '06*, page 635. ACM Press, New York, New York, USA, 2006.
- K. Deb, L. Thiele, M. Laumanns and E. Zitzler. Scalable multi-objective optimization test problems. 2002b.
- K. Deep and M. Thakur. A new crossover operator for real coded genetic algorithms. *Applied Mathematics and Computation*, 188(1):895–911, 2007a.
- K. Deep and M. Thakur. A new mutation operator for real coded genetic algorithms. *Applied Mathematics and Computation*, 193(1):211–230, 2007b.
- S. N. Deming and S. L. Morgan. Teaching the fundamentals of experimental design. *Analytica Chimica Acta*, 150:183–198, 1983.
- D. Dennett. The Baldwin effect: A crane, not a skyhook. *Evolution and learning: The Baldwin effect*, 2003.
- M. Deubel, G. von Freymann, M. Wegener, S. Pereira, K. Busch and C. M. Soukoulis. Direct laser writing of three-dimensional photonic-crystal templates for telecommunications. *Nature materials*, 3(7):444–7, 2004.
- L. Devroye and T. Wagner. The strong uniform consistency of nearest neighbor density estimates. *The Annals of Statistics*, 5(3):536–540, 1977.
- J. G. Digalakis and K. G. Margaritis. On benchmarking functions for genetic algorithms. *International Journal of Computer Mathematics*, 77(4):481–506, 2001.

- S. Dignum and R. Poli. Generalisation of the limiting distribution of program sizes in tree-based genetic programming and analysis of its effects on bloat. In *Proceedings of the 9th Annual Conference on Genetic and Evolutionary Computation*, volume 2, pages 1588–1595. ACM Press, London, 2007.
- F. Dill, W. Hornberger, P. Hauge and J. Shaw. Characterization of positive photoresist. *IEEE Transactions on electron devices*, 22(7):445–452, 1975.
- C. Dolainsky, W. Maurer and T. Waas. Evaluation of resist models for fast optical proximity correction. In *Proceedings of SPIE*, volume 3236, page 202. 1997.
- M. Dorigo. *Optimization, learning and natural algorithms*. Ph. d. thesis, Politecnico di Milano, Italy, 1992.
- C. Dürr. LisBON – LisBON is a Brokered Optimization Network. Technical report, Fraunhofer IISB, Erlangen, Germany, 2006a.
- C. Dürr. *Memetische Algorithmen in der globalen restringierten Optimierung von Lithographie-Prozessen*. Diplomarbeit, Universität Erlangen-Nürnberg, 2006b.
- C. Dürr, T. Fühner and P. Suganthan. LisBON: A framework for parallelisation and hybridisation of optimisation algorithms. In *2007 IEEE Congress on Evolutionary Computation*, pages 1717–1724. IEEE, 2007a.
- C. Dürr, T. Fühner, B. Tollkühn, A. Erdmann and G. Kókai. Memetic Algorithms: Parametric Optimization for Microlithography. In *Hybrid Evolutionary Algorithms*, pages 201–239. Springer, 2007b.
- M. Dusa, J. Quaedackers, O. Larsen, J. Meessen, E. van Der Heijden, G. Dicker, O. Wismans, P. de Haas, K. van Ingen Schenau, J. Finders and Others. Pitch doubling through dual-patterning lithography challenges in integration and litho budgets. In *Proceedings of SPIE*, volume 6520, page 65200G. SPIE, 2007.
- R. Eberhart and Y. Shi. Particle swarm optimization: developments, applications and resources. In *Proceedings of the 2001 Congress on Evolutionary Computation (IEEE Cat. No.01TH8546)*, volume 1, pages 81–86. IEEE, 2001.
- T. Ebihara, M. D. Levenson, W. Liu, J. He, W. Yeh, S. Ahn, T. Oga, M. Shen and H. M’saad. Beyond $k_1 = 0.25$ lithography: 70-nm L/S patterning using KrF scanners. In *Proceedings of SPIE*, volume 5256, page 985. SPIE, 2003.
- B. Edmonds. Meta-genetic programming: Co-evolving the operators of variation. *Elektrik*, 9(1):13–30, 2001.
- M. Ehrgott. *Multicriteria optimization*. 2 edition, 2005.
- A. Eiben. On evolutionary exploration and exploitation. *Fundamenta Informaticae*, 35(1-4):35–50, 1998.
- A. Eiben, E. Aarts and K. Van Hee. Global convergence of genetic algorithms: A Markov chain analysis. *Parallel Problem Solving from Nature*, 44(1):3–12, 1991.
- A. Eiben, R. Hinterding and Z. Michalewicz. Parameter control in evolutionary algorithms. *IEEE Transactions on Evolutionary Computation*, 3(2):124–141, 1999.
- A. Eiben, Z. Michalewicz, M. Schoenauer and J. Smith. Parameter Control in Evolutionary Algorithms. *Studies in Computational Intelligence*, 54:19–46, 2007.
- G. Eichfelder. *Parametergesteuerte Lösung nichtlinearer multikriterieller Optimierungsprobleme*. Ph.D. dissertation, University Erlangen-Nuremberg, 2006.

- A. Endo, H. Hoshino, T. Suganuma, M. Moriya, T. Ariga, Y. Ueno, M. Nakano, T. Asayama, T. Abe, H. Komori and Others. Laser produced EUV light source development for HVM. In *Proc. SPIE*, volume 6517, page 65170O. 2007.
- A. P. Engelbrecht. *Computational Intelligence*. John Wiley & Sons, Ltd, Chichester, UK, 2nd edition, 2007.
- A. Erdmann. Topography effects and wave aberrations in advanced PSM technology. In *Proceedings of SPIE*, volume 4346, pages 345–355. SPIE, 2001.
- A. Erdmann. Mask Modeling in the Low k_1 and Ultrahigh NA Regime: Phase and Polarization Effects. In *Proceedings of SPIE*, volume 5835, pages 69–81. SPIE, 2005.
- A. Erdmann. Optical Lithography: Technology, Physical Effects, and Modelling. lecture notes, 2010.
- A. Erdmann, G. Citarella, P. Evanschitzky, H. Schermer, V. Philipsen and P. D. Bisschop. Validity of the Hopkins approximation in simulations of hyper-NA ($NA > 1$) line-space structures for an attenuated PSM mask. In *Proceedings of SPIE*, volume 6154, pages 61540G–61540G–12. SPIE, 2006a.
- A. Erdmann, P. Evanschitzky, G. Citarella, T. Fühner and P. De Bisschop. Rigorous mask modeling using waveguide and FDTD methods: an assessment for typical hyper-NA imaging problems. In *Proceedings of SPIE*, volume 6283, page 628319. 2006b.
- A. Erdmann and C. Friedrich. Rigorous diffraction analysis for future mask technology. In *Proceedings of SPIE*, volume 4000, page 684. 2000.
- A. Erdmann, T. Fühner and P. Evanschitzky. Optimization of mask absorber stacks and illumination settings for contact hole imaging. In *Proceedings of SPIE*, volume 7028, pages 70283L–70283L–12. SPIE, 2008.
- A. Erdmann, T. Fühner, S. Seifert, S. Popp and P. Evanschitzky. The impact of the mask stack and its optical parameters on the imaging performance. In D. G. Flagello (editor), *Proceedings of SPIE*, volume 6520, pages 65201I–65201I–12. SPIE, 2007.
- A. Erdmann, C. L. Henderson and C. G. Willson. Impact of exposure induced refractive index changes of photoresists on the photolithographic process. *Journal of Applied Physics*, 89(12):8163, 2001.
- A. Erdmann and W. Henke. Simulation of optical lithography. In *Proceedings of SPIE*, volume 3729, page 480. 1999.
- A. Erdmann, C. Kalus, T. Schmoeller and A. Wolter. Efficient simulation of light diffraction from three-dimensional EUV masks using field decomposition techniques. In *Proceedings of SPIE*, volume 5037, page 482. 2003a.
- A. Erdmann, T. Schmöller and P. Evanschitzky. Masks Induced Imaging Artifacts in Extreme Ultraviolet Lithography. In *2nd International SEMATECH EUVL Symposium*. 2003b.
- A. Erdmann, F. Shao, P. Evanschitzky and T. Fühner. Mask-topography-induced phase effects and wave aberrations in optical and extreme ultraviolet lithography. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 28(6):C6J1, 2010.
- L. J. Eshelman, R. Caruana and J. D. Schaffer. Biases in the crossover landscape. In *Proceedings of the third international conference on Genetic algorithms*, pages 10–19. Morgan Kaufmann Publishers Inc., 1989.
- L. J. Eshelman and J. D. Schaffer. Real-coded genetic algorithms and interval-schemata. *Foundations of genetic algorithms*, 2(1993):187–202, 1993.

- L. J. Eshelman and J. D. Schaffer. Productive recombination and propagating and preserving schemata. *Foundations of genetic algorithms*, 3:299–313, 1995.
- S. Esquivel, H. Leiva and R. Gallard. Multiplicity in genetic algorithms to face multicriteria optimization. In *Evolutionary Computation, 1999. CEC 99. Proceedings of the 1999 Congress on*, volume 1, pages 85–90. IEEE, 1999.
- A. Estroff, Y. Fan, A. Bourov, B. Smith, P. Foubert, L. Leunissen, V. Philipsen and Y. Aksenov. Mask-induced polarization effects at high NA. In *Proc. SPIE*, volume 5754, pages 555–566. Citeseer, 2005.
- P. Evanschitzky and A. Erdmann. Three dimensional EUV simulations: a new mask near field and imaging simulation system. In *Proceedings of SPIE*, volume 5992, pages 59925B–59925B–9. SPIE, 2005.
- P. Evanschitzky, A. Erdmann, M. Besacier and P. Schiavone. Simulation of extreme ultraviolet masks with defective multilayers. In *Proceedings of SPIE*, volume 5130, page 1035. 2003.
- P. Evanschitzky, A. Erdmann and T. Fühner. Extended Abbe approach for fast and accurate lithography imaging simulations. In U. F. W. Behringer (editor), *Proceedings of SPIE*, volume 7470, pages 747007–747007–11. SPIE, 2009.
- P. Evanschitzky, F. Shao, A. Erdmann and D. Reibold. Simulation of larger mask areas using the waveguide method with fast decomposition technique. In R. J. Naber and H. Kawahira (editors), *Photomask Technology 2007*, volume 6730, pages 67301P–9. SPIE, Monterey, CA, USA, 2007.
- P. Evanschitzky, F. Shao, T. Fühner and A. Erdmann. Compensation of mask induced aberrations by projector wavefront control. In M. V. Dusa (editor), *Proceedings of SPIE*, volume 7973, pages 797329–797329–8. 2011.
- S. Fahlman. Faster-learning variations on back-propagation: An empirical study. In *Proceedings of the 1988 connectionist models*, pages 38–51. 1988.
- Y. Fan, N. V. Lafferty, A. Bourov, L. V. Zavyalova and B. W. Smith. Study of air-bubble-induced light scattering effect on image quality in 193-nm immersion lithography. In B. W. Smith (editor), *Optical Microlithography XVII*, volume 5377, pages 477–486. SPIE, 2004.
- H. Fang, Q. Wang, Y. Y.-C. Tu and M. F. Horstemeyer. An efficient non-dominated sorting method for evolutionary algorithms. *Evolutionary computation*, 16(3):355–384, 2008.
- J. Farmer, N. H. Packard and A. S. Perelson. The immune system, adaptation, and machine learning. *Physica D: Nonlinear Phenomena*, 22(1-3):187–204, 1986.
- V. Farys, C. Charpin-Nicolle, M. Richard, J.-Y. Robic, V. Muffato, E. Quesnel, S. Postnikov, P. Schiavone, M. Besacier, R. Smaali and P. Naulleau. Printability of nonsmoothed buried defects in extreme ultraviolet lithography mask blanks. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 23(6):2860, 2005.
- R. Feder, E. Spiller and J. Topalian. X-ray lithography. *Polymer Engineering and Science*, 17(6):385–389, 1977.
- T. Feigl, H. Lauth, S. Yulin and N. Kaiser. Heat resistance of EUV multilayer mirrors for long-time applications. *Microelectronic Engineering*, 57-58:3–8, 2001.
- R. A. Ferguson, L. W. Liebmann, S. M. Mansfield, D. S. O’grady and A. K.-K. Wong. Exact transmission balanced alternating phase-shifting mask for photolithography. 1998.
- A. Fettig. *Twisted network programming essentials*. O’Reilly Media, 2005.

- J. R. Fienup. Reconstruction of an object from the modulus of its Fourier transform. *Optics Letters*, 3(1):27, 1978.
- J. Finders, M. Dusa, P. Nikolsky, Y. van Dommelen, R. Watso, T. Vandeweyer, J. Beckaert, B. Laenens and L. Van Look. Litho and patterning challenges for memory and logic applications at the 22-nm node. In *Proceedings of SPIE*, volume 7640, pages 76400C–76400C–10. 2010.
- D. G. Flagello. *High Numerical Aperture Imaging in Homogeneous Thin Films*. Ph.D. dissertation, The University of Arizona, 1993.
- D. G. Flagello, B. Arnold, S. Hansen, M. Dusa, R. J. Socha, J. Mulken and R. Garreis. Optical lithography in the sub-50-nm regime. *Proceedings of SPIE*, 5377:21–33, 2004.
- D. G. Flagello, J. de Klerk, G. Davies, R. Rogoff, B. Geh, M. Arnz, U. Wegmann and M. Kraemer. Toward a comprehensive control of full-field image quality in optical photolithography. In *Proceedings of SPIE*, volume 3051, page 672. SPIE, 1997.
- D. G. Flagello, B. Geh, S. Hansen and M. Totzeck. Polarization effects associated with hyper-numerical-aperture (> 1) lithography. *Journal of Microlithography, Microfabrication, and Microsystems*, 4(3):31104, 2005.
- D. G. Flagello and T. Milster. Three-dimensional modeling of high-numerical-aperture imaging in thin films. In *Proceedings of SPIE*, volume 1625, pages 246–261. 1992.
- R. Flavell. A new goal programming formulation. *Omega*, 4(6):731–732, 1976.
- R. Fletcher. A new low rank quasi-Newton update scheme for nonlinear programming. In *Proceedings of the 22nd IFIP TC7 Conference*, pages 275–293. 2006.
- R. Fletcher and M. J. D. Powell. A Rapidly Convergent Descent Method for Minimization. *The Computer Journal*, 6(2):163–168, 1963.
- A. Fleury, E. Saint-Christophe, H. Frémont, M. Fathi, G. N’kaoua and Y. Danto. Laser Beam Lithography for Direct Patterning of Interconnections on Prediffused ASIC’s. *J. Phys. III France*, 5:1455–1467, 1995.
- T. C. Fogarty. Varying the probability of mutation in the genetic algorithm. In *Proceedings of the third international conference on Genetic algorithms*, pages 104–109. 1989.
- D. Fogel. An evolutionary approach to the traveling salesman problem. *Biological Cybernetics*, 60(2):139–144, 1988.
- D. Fogel and J. Atmar. Comparing genetic operators with Gaussian mutations in simulated evolutionary processes using linear systems. *Biological Cybernetics*, 63(2):111–114, 1990.
- D. Fogel, L. Fogel and J. Atmar. Meta-evolutionary programming. In *Conference Record of the Twenty-Fifth Asilomar Conference on Signals, Systems & Computers*, pages 540–545. IEEE Comput. Soc. Press, 1991.
- D. Fogel and A. Ghoseil. A note on representations and variation operators. *IEEE Transactions on Evolutionary Computation*, 1(2):159–161, 1997.
- D. Fogel and L. Stayton. On the effectiveness of crossover in simulated evolutionary optimization. *BioSystems*, 32(3):171–182, 1994.
- L. Fogel, A. Owens and M. Walsh. *Artificial intelligence through simulated evolution*. John Wiley, 1966.
- C. M. Fonseca and P. Fleming. Genetic algorithms for multiobjective optimization: Formulation, discussion and generalization. In *Proceedings of the fifth international conference on genetic algorithms*, volume 423, pages 416–423. Citeseer, 1993.

- C. M. Fonseca, M. Somervell, S. Scheer, W. Printz, K. Nafus, S. Hatakeyama, Y. Kuwahara, T. Niwa, S. Bernard and R. Gronheid. Advances and challenges in dual-tone development process optimization. In *Proceedings of SPIE*, volume 7274, page 72740I. 2009.
- C. M. Fonseca and P. J. Fleming. On the performance assessment and comparison of stochastic multi-objective optimizers. In *Parallel Problem Solving from Nature (PPSN IV)*, pages 584–593. Springer, 1996.
- S. Forrest. Documentation for prisoner’s dilemma and norms programs that use the genetic algorithm. techreport, University of Michigan, Ann Arbor, Michigan, 1985.
- S. Forrest and M. Mitchell. Relative building-block fitness and the building-block hypothesis. *Ann Arbor*, 1001:48109, 1993a.
- S. Forrest and M. Mitchell. What makes a problem hard for a genetic algorithm? Some anomalous results and their explanation. *Machine Learning*, 13(2):285–319, 1993b.
- J. Foulkes and R. Blaikie. Finite element simulation of absorbance modulation optical lithography. In *2008 International Conference on Nanoscience and Nanotechnology*, pages 184–187. IEEE, 2008.
- B. Fox and M. McMahon. Genetic operators for sequencing problems. *Foundations of genetic algorithms*, 1:284–300, 1991.
- M. Frank and P. Wolfe. An algorithm for quadratic programming. *Naval research logistics quarterly*, 3(1-2):95–110, 1956.
- B. Freisleben and P. Merz. A Genetic Local Search Algorithm for Solving Symmetric and Asymmetric Traveling Salesman Problems. In *IEEE Congress on Evolutionary Computation CEC 96*, Fb 12, pages 616–621. IEEE Press, 1996.
- R. H. French, V. Liberman, H. V. Tran, J. Feldman, D. J. Adelman, R. C. Wheland, W. Qiu, S. J. McLain, O. Nagao, M. Kaku, M. Mocella, M. K. Yang, M. F. Lemon, L. Brubaker, A. L. Shoe, B. Fones, B. E. Fischel, K. Krohn, D. Hardy and C. Y. Chen. *High-index immersion lithography with second-generation immersion fluids to enable numerical apertures of 1.55 for cost effective 32-nm half pitches*, volume 4. SPIE, 2007.
- M. Fritze, B. Tyrrell, D. K. Astolfi, D. Yost, P. Davis, B. Wheeler, R. D. Mallen, J. Jarmolowicz, S. G. Cann, H.-Y. Liu, M. Ma, D. Y. Chan, P. D. Rhyins, C. Carney, J. E. Ferri and B. A. Blachowicz. 100-nm node lithography with KrF? In C. J. Proglar (editor), *Optical Microlithography XIV*, volume 4346, pages 191–204. SPIE, Santa Clara, CA, USA, 2001.
- M. Fritze, B. Tyrrell, T. Fedynyshyn, M. Rothschild and P. Brooker. High-throughput hybrid optical maskless lithography: all-optical 32-nm node imaging. In *Proceedings of SPIE*, volume 5751, page 1058. SPIE, 2005.
- C. Fu, T. Yang and D. Stone. Enhancement of lithographic patterns by using serif features. *IEEE Transactions on Electron Devices*, 38(12):2599–2603, 1991.
- D. Fuard, M. Besacier and P. Schiavone. Assessment of different simplified resist models. In *Proceedings of SPIE*, volume 4691, pages 1266–1277. SPIE, 2002.
- T. Fühner, A. Erdmann and P. Evanschitzky. Simulation-based EUV source and mask optimization. In *Proceedings of SPIE*, volume 7122, pages 71221Y–71221Y–14. SPIE, 2008.
- T. Fühner, A. Erdmann and S. Seifert. Direct optimization approach for lithographic process conditions. *Journal of Micro/Nanolithography, MEMS and MOEMS*, 6(3):031006, 2007a.
- T. Fühner, P. Evanschitzky and A. Erdmann. Mutual source, mask and projector pupil optimization. In *Proceedings of SPIE*, volume 8326, pages 83260I–83260I–12. 2012.

- T. Fühner and T. Jung. Use of genetic algorithms for the development and optimization of crystal growth processes. *Journal of Crystal Growth*, 266(1-3):229–238, 2004.
- T. Fühner, S. Popp and T. Jung. A novel framework for distributing computations DisPyTE – distributing Python tasks environment. *Journal of Computational Electronics*, 5(4):349–352, 2006.
- T. Fühner, T. Schnattinger, G. Ardelean and A. Erdmann. Dr.LiTHO: a development and research lithography simulator. In *Proceedings of SPIE*, page 65203F. SPIE, 2007b.
- J. Fuhrmann, A. Fiebach, A. Erdmann and P. Trefonas. Acid diffusion effects between resists in freezing processes used for contact hole patterning. *Microelectronic Engineering*, 87(5-8):951–954, 2010.
- K. Funahashi. On the approximate realization of continuous mappings by neural networks. *Neural Networks*, 2(3):183–192, 1989.
- S. I. Gallant. Perceptron-based learning algorithms. *IEEE Transactions on Neural Networks*, 1(2):179–191, 1990.
- G. Gallatin. High-numerical-aperture scalar imaging. *Applied Optics*, 40(28):4958–4964, 2001.
- H. Gamo. Matrix treatment of partial coherence. *Progress in Optics*, 31:187–187, 1963.
- R. Gämperle, S. Müller and P. Koumoutsakos. A parameter study for differential evolution. *Advances in Intelligent Systems, Fuzzy Systems, Evolutionary Computation*, 10:293–298, 2002.
- J. Gan and K. Warwick. A genetic algorithm with dynamic niche clustering for multimodal function optimisation. In *Proceedings of International Conference on Artificial Neural Network and Genetic Algorithms*, pages 248–255. Citeseer, 1999.
- J. Garofalo, C. Biddick, R. Kostelak and S. Vaidya. Mask assisted off-axis illumination technique for random logic. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 11(6):2651–2658, 1993.
- T.-S. Gau. Customized illumination aperture filter for low k. In *Proceedings of SPIE*, volume 4000, pages 271–282. SPIE, 2000.
- B. Geh, D. G. Flagello, C. Proglar and P. Martin. The impact of mask birefringence on hyper-NA (NA> 1.0) polarized imaging. *Proceedings of SPIE*, 5992:599210–599210–15, 2005.
- F. Gembicki and Y. Y. Haimes. Approach to performance and sensitivity multiobjective optimization: The goal attainment method. *IEEE Transactions on Automatic Control*, 20(6):769–771, 1975.
- A. M. Geoffrion, J. S. Dyer and A. Feinberg. An Interactive Approach for Multi-Criterion Optimization, with an Application to the Operation of an Academic Department. *Management Science*, 19(4-Part-1):357–368, 1972.
- R. W. Gerchberg and W. O. Saxton. A practical algorithm for the determination of the phase from image and diffraction plane pictures. *Optik*, 35:237–246, 1972.
- K. Ghandehari, B. LaFontaine and B. Singh. Attenuating extreme ultraviolet (EUV) phase-shifting mask fabrication method. United States Patent Number 6673524 B2, 2001.
- L. Ghislain, V. Elings, K. Crozier, S. Manalis, S. Minne, K. Wilder, G. Kino and C. Quate. Near-field photolithography with a solid immersion lens. *Applied Physics Letters*, 74:501, 1999.
- D. Gil, R. Menon and H. I. Smith. The case for diffractive optics in maskless lithography. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 21(6):2810, 2003.
- A. Gillies. *Machine learning procedures for generating image domain feature detectors*. Ph.D. dissertation, University of Michigan, 1985.

- B. Gleason and W.-H. Cheng. Optical properties of alternating phase-shifting masks. In *Proceedings of SPIE*, volume 6349, pages 63491B–63491B–10. SPIE, 2006.
- F. Glover. Future paths for integer programming and links to artificial intelligence. *Computers & Operations Research*, 13(5):533–549, 1986.
- C.-K. Goh, Y.-S. Ong and K. C. Tan (editors). *Multi-Objective Memetic Algorithms*, volume 171 of *Studies in Computational Intelligence*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009.
- K. S. Goh, A. Lim and B. Rodrigues. Sexual Selection for Genetic Algorithms. *Artif. Intell. Rev.*, 19(2):123–152, 2003.
- D. E. Goldberg, K. Deb and J. Clark. Genetic algorithms, noise, and the sizing of populations. *Complex systems*, 6(4):333–362, 1992a.
- D. E. Goldberg. Simple genetic algorithms and the minimal, deceptive problem. *Genetic algorithms and simulated annealing*, 74:88, 1987.
- D. E. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley, Reading, MA, 1989.
- D. E. Goldberg. A note on Boltzmann tournament selection for genetic algorithms and population-oriented simulated annealing. *Complex Systems*, 4(4):445–460, 1990.
- D. E. Goldberg. Real-coded genetic algorithms, virtual alphabets, and blocking. *Complex Systems*, 5(2):139–167, 1991.
- D. E. Goldberg. *The Design of Innovation: Lessons from and for Competent Genetic Algorithms*. Kluwer Academic Publishers, Norwell, MA, USA, 2002.
- D. E. Goldberg and K. Deb. A comparative analysis of selection schemes used in genetic algorithms. *Foundations of genetic algorithms*, 1:69–93, 1991.
- D. E. Goldberg, K. Deb, H. Kargupta and G. Harik. Rapid Accurate Optimization of Difficult Problems Using Fast Messy Genetic Algorithms. In *Proceedings of the 5th International Conference on Genetic Algorithms*, pages 56–64. Morgan Kaufmann Publishers Inc., 1993a.
- D. E. Goldberg, K. Deb and D. Thierens. Toward a better understanding of mixing in genetic algorithms. *Urbana*, 51:61801, 1992b.
- D. E. Goldberg, K. Deb and D. Thierens. Toward a better understanding of mixing in genetic algorithms. *Journal of the Society of Instrument and Control Engineers*, 32(1):10–16, 1993b.
- D. E. Goldberg, B. Korb and K. Deb. Messy genetic algorithms: Motivation, analysis, and first results. *Complex systems*, 3(5):493–530, 1989.
- D. E. Goldberg and R. Lingle Jr. Alleles Loci and the Traveling Salesman Problem. In *Proceedings of the 1st International Conference on Genetic Algorithms*, pages 154–159. L. Erlbaum Associates Inc., Hillsdale, NJ, USA, 1985.
- D. E. Goldberg and P. Segrest. Finite Markov chain analysis of genetic algorithms. In *Proceedings of the Second International Conference on Genetic Algorithms on Genetic algorithms and their application*, pages 1–8. L. Erlbaum Associates Inc., Hillsdale, NJ, USA, 1987.
- D. E. Goldberg and R. E. Smith. Nonstationary function optimization using genetic algorithm with dominance and diploidy. In *Proceedings of the Second International Conference on Genetic Algorithms on Genetic algorithms and their application*, pages 59–68. L. Erlbaum Associates Inc., Hillsdale, NJ, USA, 1987.

- D. E. Goldberg and L. Wang. Adaptive niching via coevolutionary sharing. In D. Quagliarella, J. Périaux, C. Poloni and G. Winter (editors), *Genetic Algorithms and Evolution Strategy in Engineering and Computer Science*, pages 21–38. Citeseer, Chichester, UK, 1998.
- K. A. Goldberg, A. Barty, Y. Liu, P. Kearney, Y. Tezuka, T. Terasawa, J. S. Taylor, H.-S. Han and O. R. Wood. Actinic inspection of extreme ultraviolet programmed multilayer defects and cross-comparison measurements. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 24(6):2824, 2006.
- D. Goldfarb. A family of variable metric methods derived by variational means. *Mathematics of Computation*, 24:23–26, 1970.
- Y.-j. Gong, X.-m. Hu, J. Zhang, O. Liu and H.-l. Liu. A linear map-based mutation scheme for real coded genetic algorithms. In *Evolutionary Computation (CEC), 2010 IEEE Congress on*, pages 1–7. IEEE, 2010.
- J. G. Goodberlet. Patterning 100 nm features using deep-ultraviolet contact photolithography. *Applied Physics Letters*, 76(6):667, 2000.
- J. W. Goodman. *Introduction to Fourier optics*. McGraw-Hill Science, Engineering & Mathematics, 1996.
- R. Gordon, C. A. Mack and J. Petersen. Design and analysis of manufacturable alternating phase-shifting masks. In *Proceedings of SPIE*, volume 3546, pages 606–616. SPIE, 1998.
- N. Gould and P. Toint. SQP methods for large-scale nonlinear programming. *System Modelling and Optimization: Methods, Theory and Applications*, pages 149–178, 2000.
- J. Gramss, A. Stoeckel, U. Weidenmueller, H. Doering, M. Bloecker, M. Sczyrba, M. Finken, T. Wandel and D. Melzer. Multi-shaped e-beam technology for mask writing. In *Proceedings of SPIE*, volume 7823, page 782309. 2010.
- G. Granet and B. Guizal. Efficient implementation of the coupled-wave method for metallic lamellar gratings in TM polarization. *Journal of the Optical Society of America A*, 13(5):1019, 1996.
- Y. Granik. Source optimization for image fidelity and throughput. *Journal of Microlithography, Microfabrication, and Microsystems*, 3(4):509, 2004.
- Y. Granik. Fast pixel-based mask optimization for inverse lithography. *Journal of Microlithography, Microfabrication, and Microsystems*, 5:43002, 2006.
- Y. Granik and N. Cobb. MEEF as a matrix. In G. T. Dao and B. J. Grenon (editors), *Proceedings of SPIE*, volume 4562, pages 980–991. SPIE, Monterey, CA, USA, 2002.
- Y. Granik, K. Sakajiri and S. Shang. On objectives and algorithms of inverse methods in microlithography. In *Proceedings of SPIE*, volume 6349, pages 63494R–63494R–8. SPIE, 2006.
- F. Gray. Pulse code in communication. United States Patent Number 2632058, 1953.
- S. Greco, B. Matarazzo and R. Slowinski. Interactive Evolutionary Multiobjective Optimization using Dominance-based Rough Set Approach. In *IEEE Congress on Evolutionary Computation*, pages 1–8. IEEE, 2010.
- F. Greene. A method for utilizing diploid/dominance in genetic search. In *Proceedings of the First IEEE Conference on Evolutionary Computation. IEEE World Congress on Computational Intelligence*, pages 439–444. IEEE, 1994.

- R. T. Greenway, K. Jeong, A. B. Kahng, C.-H. Park and J. S. Petersen. 32nm 1-D regular pitch SRAM bitcell design for interference-assisted lithography. In H. Kawahira and L. S. Zurbrick (editors), *Proceedings of SPIE*, pages 71221L–71221L–12. SPIE, 2008.
- G. Greenwood and X. Hu. Fitness functions for multiple objective optimization problems: Combining preferences with Pareto rankings. *Foundations of genetic algorithms*, 4:437–455, 1997.
- J. J. Grefenstette and R. Gopal. Genetic algorithms for the traveling salesman problem. *Proceedings of the 1st International Conference on Genetic Algorithms*, pages 160–168, 1985.
- J. J. Grefenstette. Optimization of control parameters for genetic algorithms. *IEEE Transactions on Systems, Man and Cybernetics*, 16(1):122–128, 1986.
- J. J. Grefenstette. Deception considered harmful. *Foundations of Genetic Algorithms 2*, pages 75–91, 1993.
- J. J. Grefenstette. Virtual genetic algorithms: First Results. Technical report, Navy Center for Applied Research, Washington, D.C., 1995.
- J. J. Grefenstette and J. E. Baker. How Genetic Algorithms Work: A Critical Look at Implicit Parallelism. In *Proceedings of the 3rd International Conference on Genetic Algorithms*, pages 20–27. Morgan Kaufmann Publishers Inc., 1989.
- A. O. Griewank. Generalized descent for global optimization. *Journal of Optimization Theory and Applications*, 34(1):11–39, 1981.
- W. Gropp, E. Lusk and A. Skjellum. *Using MPI: Portable Parallel Programming with the Message-Passing Interface*. MIT Press, 1999.
- P. B. Grosso. *Computer simulations of genetic adaptation: Parallel Subcomponent Interaction in a Multilocus Model*. Ph.D. dissertation, University of Michigan, Ann Arbor, MI, 1985.
- V. Grunert da Fonseca, C. M. Fonseca and A. O. Hall. Inferential performance assessment of stochastic optimisers and the attainment function. *Evolutionary Multi-Criterion Optimization*, 1993:213–225, 2001.
- T. Grüninger and D. Wallace. Multimodal optimization using genetic algorithms. Technical report, Massachusetts Institute of Technology, 1996.
- A. Gu and A. Zakhor. Optical Proximity Correction With Linear Regression. *IEEE Transactions on Semiconductor Manufacturing*, 21(2):263–271, 2008.
- X. Gu, A. J. Berro, Y. Cho, K. Jen, S. Lee, T. Ngai, T. Ogata, W. J. Durand, A. Sundaresan, J. R. Lancaster, S. Jockusch, P. Zimmerman, N. J. Turro and C. G. Willson. Fundamental study of optical threshold layer approach towards double exposure lithography. In *Proceedings of SPIE*, volume 7273, pages 72731C–72731C–11. SPIE, 2009.
- D. J. Guerrero, S. Gibbons, J. Lowes and R. Mercado. Anti-reflective coating for multipatterning lithography. In *Proceedings of SPIE*, volume 6923, pages 69230X–69230X–7. SPIE, 2008.
- R. Guerrieri, K. Tadros, J. Gamelin and A. Neureuther. Massively parallel algorithms for scattering in optical lithography. *IEEE Trans. Computer-Aided Design*, 10(9):1091–1100, 1991.
- F. G. Guimarães, F. Campelo, H. Igarashi, D. A. Lowther and J. A. Ramírez. Optimization of Cost Functions Using Evolutionary Algorithms With Local Learning and Local Search. *IEEE Transactions on Magnetics*, 43(4):1641–1644, 2007.

- F. G. Guimarães, E. Wanner, F. Campelo, R. Takahashi, H. Igarashi, D. Lowther and J. Ramirez. Local Learning and Search in Memetic Algorithms. In *2006 IEEE International Conference on Evolutionary Computation*, pages 2936–2943. IEEE, 2006.
- E. Gullikson, C. Cerjan, D. G. Stearns, P. Mirkarimi and D. Sweeney. Practical approach for modeling extreme ultraviolet lithography mask defects. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 20:81, 2002.
- S. Gustafson and E. K. Burke. The Speciating Island Model: An alternative parallel evolutionary algorithm. *Journal of Parallel and Distributed Computing*, 66(8):1025–1036, 2006.
- H. Hada, T. Watanabe, K. Hamamoto, H. Kinoshita and H. Komano. Evaluation of resist outgassing by EUV irradiation. In *Proceedings of SPIE*, volume 5374, page 686. 2004.
- C. Hafner. *The Generalized Multipole Technique for Computational Electromagnetics (Antenna library)*. Artech House, 1990.
- R. Haight, D. Hayden, P. Longo, T. Neary and A. Wagner. Implementation and performance of a femtosecond laser mask repair system in manufacturing. In *Proceedings of SPIE*, volume 3546, page 477. SPIE, 1998.
- Y. Y. Haimes. Modeling and control of the pollution of water resources systems via multilevel approach. *Journal of the American Water Resources Association*, 7(1):93–101, 1971.
- P. Hajela and C. Lin. Genetic search strategies in multicriterion optimal design. *Structural and Multidisciplinary Optimization*, 4(2):99–107, 1992.
- K. Hamamoto, Y. Tanaka, S. Y. Lee, N. Hosokawa, N. Sakaya, M. Hosoya, T. Shoki, T. Watanabe and H. Kinoshita. Mask defect inspection using an extreme ultraviolet microscope. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 23(6):2852, 2005.
- L. Han and M. Neumann. Effect of dimensionality on the Nelder–Mead simplex method. *Optimization Methods and Software*, 21(1):1–16, 2006.
- S. P. Han. Superlinearly convergent variable metric algorithms for general nonlinear programming problems. *Mathematical Programming*, 11(1):263–282, 1976.
- S. P. Han. A globally convergent method for nonlinear programming. *Journal of optimization theory and applications*, 22(3):297–309, 1977.
- H. Hänsel and W. Polack. Verfahren zur Herstellung einer Phasenmaske mit Amplitudenstruktur. 1976.
- A. P. Hansen and M. Jaszkiwicz. Evaluating the quality of approximations to the non-dominated set. Technical report, Technical University of Denmark, Lyngby, Denmark, 1998.
- R. M. Haralick, S. R. Sternberg and X. Zhuang. Image analysis using mathematical morphology. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 9(4):532–550, 1987.
- G. Harik. Finding multimodal solutions using restricted tournament selection. In L. Eshelman (editor), *Proceedings of the Sixth International Conference on Genetic Algorithms*, pages 24–31. Citeseer, San Francisco, CA, 1995.
- G. Harik. *Learning linkage to efficiently solve problems of bounded difficulty using genetic algorithms*. Ph.D. dissertation, University of Michigan, Ann Arbor, Michigan, 1997.
- G. Harik, E. Cantú-Paz, D. E. Goldberg and B. L. Miller. The Gambler’s Ruin Problem, Genetic Algorithms, and the Sizing of Populations. *Evolutionary Computation*, 7(3):231–253, 1999.

- G. Harik and F. G. Lobo. A parameter-less genetic algorithm. In *Proceedings of the Genetic and Evolutionary Computation Conference*, volume 1, pages 258–265. Citeseer, 1999.
- S. S. Harilal, B. O'Shay, Y. Tao and M. Tillack. Ion debris mitigation from tin plasma using ambient gas, magnetic field and combined effects. *Applied Physics B*, 86(3):547–553, 2006.
- S. S. Harilal, B. O'Shay and M. S. Tillack. Debris mitigation in a laser-produced tin plume using a magnetic field. *Journal of Applied Physics*, 98(3):036102, 2005.
- W. E. Hart and R. Belew. Optimizing an arbitrary function is hard for the genetic algorithm. In *Proceedings of the Fourth International Conference on Genetic Algorithms*, pages 190–195. 1991.
- W. E. Hart. *Adaptive Global Optimization with Local Search*. Ph.D. dissertation, Univ. California, San Diego, CA., 1994.
- W. E. Hart, J. E. Smith and N. Krasnogor (editors). *Recent Advances in Memetic Algorithms*, volume 166 of *Studies in Fuzziness and Soft Computing*. Springer-Verlag, Berlin/Heidelberg, 2005.
- M. H. Hassoun. *Fundamentals of artificial neural networks*, volume 29. The MIT Press, 1995.
- E. Hecht. *Optics*. Addison-Wesley, 2002.
- R. Heim. Practical aspects of contact/proximity, photomask/wafer exposure. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 100, page 104. 1977.
- E. Hendrickx, A. Tritchkov, K. Sakajiri, Y. Granik, M. Kempsell and G. Vandenberghe. Hyper-NA imaging of 45nm node random CH layouts using inverse lithography. In *Proceedings of SPIE*, volume 6924, pages 69240L–69240L–10. SPIE, 2008.
- W. Henke, D. Mewes, M. Weiß, G. Czech and R. Schieblhoyler. Simulation of defects in 3-dimensional resist profiles in optical lithography. *Microelectronic Engineering*, 13(1-4):497–501, 1991.
- W. Henke, R. Schwalm, M. Weiß and J. Pelka. Diffraction effects in submicron contact or proximity printing. *Microelectronic Engineering*, 10(2):73–89, 1990.
- W. Henke and M. Torkler. Modeling of edge roughness in ion projection lithography. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 17(6):3112, 1999.
- F. Herrera. Tuning fuzzy logic controllers by genetic algorithms. *International Journal of Approximate Reasoning*, 12(3-4):299–315, 1995.
- F. Herrera, M. Lozano and A. Sánchez. A taxonomy for the crossover operator for real-coded genetic algorithms: An experimental study. *International Journal of Intelligent Systems*, 18(3):309–338, 2003.
- F. Herrera, M. Lozano and J. Verdegay. Tackling real-coded genetic algorithms: Operators and tools for behavioural analysis. *Artificial Intelligence Review*, 12(4):265–319, 1998.
- A. Hertz. Tabu search for large scale timetabling problems. *European Journal of Operational Research*, 1991.
- J. Hesser and R. Männer. Towards an optimal mutation probability for genetic algorithms. In *Parallel problem solving from nature*, pages 23–32. Springer, 1991.
- T. Higuchi, S. Tsutsui and M. Yamamura. Theoretical analysis of simplex crossover for real-coded genetic algorithms. In M. Schoenauer, K. Deb, G. Rudolph, X. Yao, E. Lutton, J. J. Merelo and H.-P. Schwefel (editors), *Parallel Problem Solving from Nature PPSN VI*, volume 1917 of *Lecture Notes in Computer Science*, pages 365–374. Springer, Berlin, Heidelberg, 2000.

- M. D. Himel, R. E. Hutchins, J. C. Colvin, M. K. Poutous, A. D. Kathman and A. S. Fedor. Design and fabrication of customized illumination patterns for low-k₁ lithography: a diffractive approach. In C. J. Proglar (editor), *Optical Microlithography XIV*, volume 4346, pages 1436–1442. SPIE, Santa Clara, CA, USA, 2001.
- W. D. Hinsberg, F. Houle, J. Hoffnagle, M. Sanchez, G. Wallraff, M. Morrison and S. Frank. Deep-ultraviolet interferometric lithography as a tool for assessment of chemically amplified photoresist performance. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 16(6):3689, 1998.
- W. D. Hinsberg, S. A. MacDonald, L. A. Pederson and C. G. Willson. A lithographic analog of color photography: self-aligning photolithography using a resist with wavelength-dependent tone. *Journal of imaging science*, 33(4):129–135, 1989.
- R. Hinterding. Self-adaptation using multi-chromosomes. In *Proceedings of 1997 IEEE International Conference on Evolutionary Computation (ICEC '97)*, pages 87–91. IEEE, 1997.
- R. Hinterding, Z. Michalewicz and T. Peachey. Self-adaptive genetic algorithm for numeric functions. pages 420–429, 1996.
- G. E. Hinton. How learning can guide evolution. *Complex systems*, 1(3):495–502, 1987.
- Y. Hirai, S. Tomida, K. Ikeda, M. Sasago, M. Endo, S. Hayama and N. Nomura. Three-dimensional resist process simulator PEACE (photo and electron beam lithography analyzing computer engineering system). *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 10(6):802–807, 1991.
- T. Hiroyasu, S. Nakayama and M. Miki. Comparison study of SPEA2+, SPEA2, and NSGA-II in diesel engine emissions and fuel economy problem. In *Evolutionary Computation, 2005. The 2005 IEEE Congress on*, volume 1, pages 236–242. IEEE, 2005.
- C. Hocaoglu and A. Sanderson. Multimodal function optimization using minimal representation size clustering and its application to planning multipaths. *Evolutionary Computation*, 5(1):81–104, 1997.
- J. A. Hoffnagle, W. D. Hinsberg, M. Sanchez and F. A. Houle. Liquid immersion deep-ultraviolet interferometric lithography. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 17(6):3306, 1999.
- J. Holland. *Adaptation in natural and artificial systems*. University of Michigan Press, Ann Arbor, Michigan, 1975.
- R. Hollstien. *Artificial genetic adaptation in computer control systems*. Ph.D. dissertation, University of Michigan, 1971.
- A. Homaifar, S. Guan and G. E. Liepins. A New Approach on the Traveling Salesman Problem by Genetic Algorithms. In *Proceedings of the 5th International Conference on Genetic Algorithms*, pages 460–466. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1993.
- Y.-S. Hong, H. Lee and M.-J. Tahk. Acceleration of the convergence speed of evolutionary algorithms using multi-layer neural networks. *Engineering Optimization*, 35(1):91–102, 2003.
- H. H. Hopkins. The Concept of Partial Coherence in Optics. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 208(1093):263–277, 1951.
- H. H. Hopkins. On the Diffraction Theory of Optical Images. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 217(1130):408–432, 1953.

- M. Hori, T. Nagai, A. Nakamura, T. Abe, G. Wakamatsu, T. Kakizawa, Y. Anno, M. Sugiura, S. Kusumoto, Y. Yamaguchi and Others. Sub-40-nm half-pitch double patterning with resist freezing process. In *Proceedings of SPIE*, volume 6923, page 69230H. SPIE, 2008.
- J. Horn, D. E. Goldberg and K. Deb. Implicit niching in a learning classifier system: Nature's way. *Evolutionary Computation*, 2(1):37–66, 1994.
- K. Hornik, M. Stinchcombe and H. White. Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5):359–366, 1989.
- W. Huang, C. Lai, B. Luo, C. Tsai, M. Chih, C. Lai, C. Kuo, R. Liu and H. Lin. Intelligent model-based OPC. In *Proceedings of SPIE*, volume 6154, page 615436. SPIE, 2006.
- J. P. Hugonin, M. Besbes and P. Lalanne. Hybridization of electromagnetic numerical methods through the G-matrix algorithm. *Optics Letters*, 33(14):1590–1592, 2008.
- S. Huh, L. Ren, D. Chan, S. Wurm, K. Goldberg, I. Mochi, T. Nakajima, M. Kishimoto, B. Ahn, I. Kang, J.-o. Park, K. Cho, S.-i. Han and T. Laursen. A study of defects on EUV masks using blank inspection, patterned mask inspection, and wafer inspection. In *Proceedings of SPIE*, pages 76360K–76360K–7. 2010.
- D. G. Humphrey and J. R. Wilson. A Revised Simplex Search Procedure for Stochastic Simulation Response Surface Optimization. *INFORMS Journal on Computing*, 12(4):272–283, 2000.
- C.-Y. Hung, B. Zhang, E. Guo, L. Pang, Y. Liu, K. Wang and G. Dai. Pushing the lithography limit: applying inverse lithography technology (ILT) at the 65nm generation. In D. G. Flagello (editor), *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 6154 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pages 562–571. 2006.
- M. Hüsken, Y. Jin and B. Sendhoff. Structure Optimization of Neural Networks for Evolutionary Design Optimization. *Soft Computing*, 9(1):21–28, 2002.
- J. Iba, K. Hashimoto, R. A. Ferguson, T. Yanagisawa and D. J. Samuels. Electrical characterization of across-field lithographic performance for 256-Mbit DRAM technologies. In H. Yoshihara (editor), *Proceedings of SPIE*, volume 2512, pages 218–225. SPIE, Kanagawa, Japan, 1995.
- Y. Ichikawa and Y. Ishii. Retaining diversity of genetic algorithms for multivariable optimization and neural network learning. In *IEEE International Conference on Neural Networks*, pages 1110–1114. IEEE, 1993.
- C. Igel, N. Hansen and S. Roth. Covariance matrix adaptation for multi-objective optimization. *Evolutionary computation*, 15(1):1–28, 2007.
- C. Igel and M. Hüsken. Improving the Rprop learning algorithm. In *Proceedings of the second international ICSC symposium on neural computation (NC 2000)*, pages 115–121. 2000.
- M. T. Iglesias, B. Naudts, A. Verschoren and C. Vidal. *Foundations of Generic Optimization; Volume 1: A Combinatorial Approach to Epistasis*, volume 20 of *Mathematical Modelling: Theory and Applications*. Springer-Verlag, Berlin/Heidelberg, 2005.
- M. T. Iglesias, V. S. Peñaranda, C. Vidal and A. Verschoren. Higher Epistasis in Genetic Algorithms. *Bulletin of the Australian Mathematical Society*, 77(02), 2008.
- J. Ignizio. *Goal programming and extensions*. Lexington Books, Lexington, MA, 1976.
- M. Ito, T. Ogawa, K. Otaki, I. Nishiyama, S. Okazaki and T. Terasawa. Simulation of multilayer defects in extreme ultraviolet masks. *Japanese Journal of Applied Physics*, 40:2549, 2001.
- J. D. Jackson. *Classical Electrodynamics*. Wiley & Sons, New York, NY, USA, 1962.

- C. Jacob. *Illustrating evolutionary computation with Mathematica*. Morgan Kaufmann, 2001.
- J. Jahn. *Vector Optimization*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2 edition, 2011.
- C. Z. Janikow and Z. Michalewicz. A specialized genetic algorithm for numerical optimization problems. In *Tools for Artificial Intelligence, 1990., Proceedings of the 2nd International IEEE Conference on*, pages 798–804. IEEE, 1990.
- C. Z. Janikow and Z. Michalewicz. An experimental comparison of binary and floating point representations in genetic algorithms. In *Proceedings of the Fourth International Conference on Genetic Algorithms*, pages 31–36. San Mateo: Morgan Kaufmann Publishers, 1991.
- H. Jasper, T. Modderman, M. van De Kerkhof, C. Wagner, J. Mulken, W. de Boeij, E. van Setten and B. Kneer. Immersion lithography with an ultrahigh-NA in-line catadioptric lens and a high-transmission flexible polarization illumination system. In D. G. Flagello (editor), *Optical Microlithography XIX*, volume 6154, pages 61541W–14. SPIE, San Jose, CA, USA, 2006.
- A. Jaszkiewicz. Interactive multiobjective optimization with the Pareto memetic algorithm. *Foundations of Computing and Decision Sciences*, 32(1):15–32, 2007.
- A. Jaszkiewicz and R. Słowiński. The ‘Light Beam Search’ approach – an overview of methodology applications. *European Journal of Operational Research*, 113(2):300–314, 1999.
- M. Jensen. Reducing the Run-Time Complexity of Multiobjective EAs: The NSGA-II and Other Algorithms. *IEEE Transactions on Evolutionary Computation*, 7(5):503–515, 2003.
- N. Jia and E. Y. Lam. Machine learning for inverse lithography: using stochastic gradient descent for robust photomask synthesis. *Journal of Optics*, 12(4):045601, 2010.
- N. Jia and E. Y. Lam. Pixelated source mask optimization for process robustness in optical lithography. *Optics Express*, 19(20):19384, 2011.
- Y. Jin. A comprehensive survey of fitness approximation in evolutionary computation. *Soft Computing*, 9(1):3–12, 2005.
- Y. Jin and J. Branke. Evolutionary optimization in uncertain environments-a survey. *Evolutionary Computation, IEEE Transactions on*, 9(3):303–317, 2005.
- Y. Jin and M. Olhofer. Managing approximate models in evolutionary aerodynamic design optimization. In *Proceedings of 2001 Conference on Evolutionary Computation*, volume 1, pages 592–599. 2001.
- Y. Jin, M. Olhofer and B. Sendhoff. On evolutionary optimization with approximate fitness functions. In D. Whitley, D. Goldberg, E. Cantú-Paz, L. Spector, I. Parmee and H. G. Beyer (editors), *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 786–793. Morgan Kaufmann, 2000.
- P. Jog, J. Y. Suh and D. Van Gucht. The Effects of Population Size, Heuristic Crossover and Local Improvement on a Genetic Algorithm for the Traveling Salesman Problem. In J. D. Schaffer (editor), *Proceedings of the Third International Conference on Genetic Algorithms*, pages 110–115. Morgan Kaufman, 1989.
- R. Jonckheere, K. Ronse, O. Popa and L. Van Den Hove. Molybdenum silicide based attenuated phase-shift masks. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 12(6):3765, 1994.
- R. Jonckheere, A. K.-K. Wong, A. Yen, K. Ronse and L. Van Den Hove. Optical proximity correction: Mask pattern-generation challenges. *Microelectronic Engineering*, 30(1-4):115–118, 1996.

- K. Jones, I. Lustig and A. Kornhasuer. Optimization techniques applied to neural networks: Line search implementation for back propagation. In *Neural Networks, 1990., 1990 IJCNN International Joint Conference on*, pages 933–939. IEEE, 1990.
- R. C. Jones. A New Calculus for the Treatment of Optical Systems. *Journal of the Optical Society of America*, 31(7):488–493, 1941.
- A. Jouve, J. Simon, J. Foucher, T. David, J. Tortai and H. H. Solak. Overcoming pattern collapse of ultra high resolution dense lines obtained with EUV resists. In *Proceedings of SPIE*, volume 5753, page 720. SPIE, 2005.
- B. A. Julstrom. What have you done for me lately? Adapting operator probabilities in a steady-state genetic algorithm. In *Proceedings of the 6th International Conference on Genetic Algorithms*, pages 81–87. 1995.
- H. Y. Jung, Y. R. Park, H. J. Lee, N.-E. Lee, C. Y. Jeong and J. Ahn. Selective dry etching of attenuated phase-shift mask materials for extreme ultraviolet lithography using inductively coupled plasmas. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 27(6):2361, 2009.
- W. Jung, C. Kim, J. Eom, S. Cho, S. Jeon, J. Kim, J. Moon, B. Lee and S. Park. Patterning with spacer for expanding the resolution limit of current lithography tool. In *Proceedings of SPIE*, volume 6156, page 61561J. SPIE, 2006.
- A. B. Kahng. Exploiting synergies of multiple crossovers: initial studies. In *Proceedings of 1995 IEEE International Conference on Evolutionary Computation*, volume 1, page 245. IEEE, 1995.
- A. B. Kahng and Y. C. Pati. Subwavelength lithography and its potential impact on design and EDA. In *Proceedings of the 36th ACM/IEEE conference on Design automation conference - DAC '99*, pages 799–804. ACM Press, New York, New York, USA, 1999.
- I. Kaliszewski. A modified weighted Tchebycheff metric for multiple objective programming. *Computers & operations research*, 14(4):315–323, 1987.
- F. Kalk, R. French, H. Alpay and G. Hughes. Attenuated phase-shifting photomasks fabricated from Cr-based embedded shifter blanks. In *Proc. SPIE*, volume 2254, page 64. 1994.
- V. Kamat, K. Green, S. Chheda, S. Muehle, V. Kolagunta, B. Wilkinson and C. Philbin. Model-based OPC methodology for 0.13 micron technology. In *Proceedings of SPIE*, volume 4562, page 727. 2002.
- K. Kamon, T. Miyamoto, Y. Myoi, H. Nagata, M. Tanaka and K. Horie. Photolithography System Using Annular Illumination. *Japanese Journal of Applied Physics*, 30(Part 1, No. 11B):3021–3029, 1991.
- H. Kang, S. Hansen, J. van Schoot and K. van Ingen Schenau. EUV simulation extension study for mask shadowing effect and its correction. In *Proceedings of SPIE*, volume 6921, page 69213I. 2008.
- H. Y. Kang and C. K. Hwangbo. Optical design of hybrid-type attenuated phase-shift masks for extreme-ultraviolet lithography by use of a Fabry-Perot interference filter. *Applied Optics*, 47(13):C75, 2007.
- D. Karaboga and B. Basturk. A powerful and efficient algorithm for numerical function optimization: artificial bee colony (ABC) algorithm. *Journal of Global Optimization*, 39(3):459–471, 2007.
- H. Kargupta. The gene expression messy genetic algorithm. In *Proceedings of the 1996 IEEE International Conference on Evolutionary Computation.*, pages 631–636. 1996.
- G. Karniadakis. Spectral element simulations of laminar and turbulent flows in complex geometries. *Applied Numerical Mathematics*, 6(1-2):85–105, 1989.

- A. Kawai, T. Niiyama, H. Endo, M. Yamanaka, A. Ishikawa, K. Suzuki, O. Tamada and M. Sanada. Adhesion and removal of micro bubbles for immersion lithography. In Q. Lin (editor), *Advances in Resist Technology and Processing XXIII*, volume 6153, page 61531S. SPIE, 2006.
- V. Kecman. *Learning and Soft Computing: Support Vector Machines, Neural Networks, and Fuzzy Logic Models*. MIT Press, Cambridge, MA, 2001.
- T. Keelin, III. *A Protocol and Procedure for Assessing Multiattribute Preference Functions*. Ph.D. dissertation, Stanford University, 1976.
- R. Keeney. *Multidimensional utility functions: theory, assessment, and application*. Phd thesis, Massachusetts Institute of Technology, 1969.
- C. T. Kelley. Detection and Remediation of Stagnation in the Nelder–Mead Algorithm Using a Sufficient Decrease Condition. *SIAM Journal on Optimization*, 10(1):43, 1999.
- G. Kendall, E. Soubeiga and P. Cowling. Choice function and random hyperheuristics. In *Proceedings of the fourth Asia-Pacific Conference on Simulated Evolution And Learning, SEAL*, pages 667–671. 2002.
- J. Kennedy and R. Eberhart. Particle swarm optimization. In *Proceedings of ICNN'95 - International Conference on Neural Networks*, volume 4, pages 1942–1948. IEEE, 1995.
- V. Khare, X. Yao and K. Deb. Performance scaling of multi-objective evolutionary algorithms. In *Evolutionary Multi-Criterion Optimization*, pages 72–72. Springer, 2003.
- P. Kik, S. Maier and H. Atwater. Plasmon printing-a new approach to near-field lithography. In *MATERIALS RESEARCH SOCIETY SYMPOSIUM PROCEEDINGS*, volume 705, pages 101–106. Warrendale, Pa.; Materials Research Society; 1999, 2001.
- D. Kim, W. Oldham and A. Neureuther. Development of positive photoresist. *IEEE Transactions on Electron Devices*, 31(12):1730–1736, 1984.
- H. Kim, D. Nam, C. Hwang, Y. Kang and S. Woo. Layer-specific illumination optimization by Monte Carlo method. *Proceedings of SPIE*, 5040:244–250, 2003.
- J. Kim, K. Park, Y. Kim, Y.-H. Oh, J. Lee, C. Go and S. Lim. Monte-carlo based optical proximity correction of DRAM storage capacitors and hammer heads. *Journal of the Korean Physical Society*, 40(1):163–166, 2002.
- R.-h. Kim and H. J. Levinson. Application of contrast enhancement layer to 193 nm lithography. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 25(6):2466, 2007.
- H. Kirchauer. *Photolithography Simulation*. Ph.D. dissertation, Technische Universität Wien, 1998.
- S. Kirkpatrick, C. D. Gelatt and M. P. Vecchi. Optimization by Simulated Annealing. *Science (New York, N.Y.)*, 220(4598):671–680, 1983.
- N. Kishi, N. Iizuka, F. Jiang, S. Orishimo, T. Kawamura, K. Horioka, M. Watanabe, A. Okino and E. Hotta. Development of Xe and Sn discharge EUV light source. *Journal of Physics: Conference Series*, 112(4):042063, 2008.
- Y. Kitahata, D. Adachi, Y. Inazuki, T. Sutou, T. Nagai, N. Toyama, T. Yokoyama, Y. Morikawa, H. Mohri and N. Hayashi. Study of Mask Material and Topography for 45-nm Node. In *3rd International Symposium on Immersion Lithography*. Kyoto, Japan, 2006.

- C. Klein, J. Klikovits, L. Szikszai, E. Platzgummer and H. Loeschner. 50 keV electron-beam projection maskless lithography (PML2): results obtained with 2,500 programmable 12.5-nm sized beams. In *Proceedings of SPIE*, volume 7637, page 76370B. 2010.
- M. Klosner and W. Silvast. Intense xenon capillary discharge extreme-ultraviolet source in the 10–16-nm-wavelength region. *optics letters*, 23(20):1609–1611, 1998.
- J. D. Knowles and D. Corne. On metrics for comparing nondominated sets. In *Proceedings of the 2002 Congress on Evolutionary Computation. CEC'02 (Cat. No.02TH8600)*, volume 1, pages 711–716. IEEE, 2002.
- J. D. Knowles and D. W. Corne. Approximating the nondominated front using the pareto archived evolution strategy. *Evolutionary computation*, 8(2):149–172, 2000.
- J. D. Knowles, L. Thiele and E. Zitzler. A tutorial on the performance assessment of stochastic multiobjective optimizers. Technical report, Computer Engineering and Networks Laboratory (TIK), Swiss Federal Institute of Technology (ETH) Zurich, 2006.
- C. Koay, C. Keyser, K. Takenoshita, E. Fujiwara, M. Al-Rabban, M. Richardson, I. Turcu, H. Rieger, A. Stone and J. Morris. High conversion efficiency tin material laser plasma source for EUVL. In *Proc. SPIE*, volume 5037, pages 801–806. 2003.
- S. Kobayashi, J. Santillan and T. Itani. Quantitative analysis of EUV resist outgassing. In *Proceedings of SPIE*, volume 6923, page 692345. SPIE, 2008.
- M. Kocsis, P. D. Bisschop, M. Maenhoudt, Y.-C. Kim, G. Wells, S. Light and T. DiBiase. A methodology for the characterization of topography induced immersion bubble defects. In B. W. Smith (editor), *Optical Microlithography XVIII*, volume 5754, pages 154–163. SPIE, 2004.
- A. Köhler. Ein neues Beleuchtungsverfahren für mikrophotographische Zwecke. *Zeitschrift für wissenschaftliche Mikroskopie und für mikroskopische Technik*, 10(4):433–440, 1893.
- T. G. Kolda, R. M. Lewis and V. J. Torczon. Optimization by direct search: New perspectives on some classical and modern methods. *SIAM review*, 45(3):385–482, 2003.
- A. N. Kolmogorov. On the Representation of continuous functions of several variables by superposition of continuous functions of one variable and addition. *Dokl. Akad. Nauk SSSR*, 114:953–956, 1957.
- I. Korejo, S. Yang and C. Li. A Directed Mutation Operator for Real Coded Genetic Algorithms. *Applications of Evolutionary Computation*, pages 491–500, 2010.
- V. Koronkevich, V. Kiriyanov, F. Kokoulin, I. Palchikova, A. Poleshchuk, A. Sedukhin, E. Churin, A. Shcherbachenko and Y. Yurlov. Fabrication of kinoform optical elements. *Optik*, 67(3):257–266, 1984.
- M. Kotera, K. Yagura, H. Tanaka, D. Kawano and T. Maekawa. Extreme Ultraviolet Lithography Simulation by Tracing Photoelectron Trajectories in Resist. *Japanese Journal of Applied Physics*, 47(No. 6):4944–4949, 2008.
- J. Koza. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*, volume 229. The MIT Press, 1992.
- J. Koza. *Genetic Programming II: Automatic Discovery of Reusable Programs*. The MIT Press, 1994.
- T. Kozawa, S. Tagawa, H. B. Cao, H. Deng and M. J. Leeson. Acid distribution in chemically amplified extreme ultraviolet resist. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 25(6):2481, 2007.

- T. Kozawa, S. Tagawa, H. Oizumi and I. Nishiyama. Acid generation efficiency in a model system of chemically amplified extreme ultraviolet resist. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 24(6):L27, 2006.
- N. Krasnogor. *Studies on the theory and design space of memetic algorithms*. Ph.D. dissertation, University of the West of England, Bristol, 2002.
- N. Krasnogor and S. Gustafson. A Study on the use of “self-generation” in memetic algorithms. *Natural Computing*, 3(1):53–76, 2004.
- N. Krasnogor and J. Smith. MAFRA: A java memetic algorithms framework. *Proceedings of the 2000 International Genetic and Evolutionary Computation Conference GECCO 2000*, pages 125–130, 2000.
- C. Krautschik, M. Ito, I. Nishiyama and T. Mori. Quantifying EUV imaging tolerances for the 70-, 50-, 35-nm modes through rigorous aerial image simulations. In *Proceedings of SPIE*, volume 4343, page 524. 2001.
- H. W. Kuhn and A. W. Tucker. Nonlinear programming. In J. Neyman (editor), *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, x, pages 481–492. Berkeley, U. of Calif. Press, 1951.
- J. C. Lagarias, B. Poonen and M. H. Wright. Convergence of the Restricted Nelder-Mead Algorithm in Two Dimensions. *SIAM Journal on Optimization*, 22(2):501–532, 2012.
- J. C. Lagarias, J. A. Reeds, M. H. Wright and P. E. Wright. Convergence Properties of the Nelder-Mead Simplex Method in Low Dimensions. *SIAM Journal on Optimization*, 9(1):112, 1998.
- P. Lalanne and G. M. Morris. Highly improved convergence of the coupled-wave method for TM polarization. *Journal of the Optical Society of America A*, 13(4):779, 1996.
- I. Lalovic. Investigation of cross-field wavefront aberrations of KrF lithography exposure systems as a function of excimer laser bandwidth. *Proceedings of SPIE*, 4346:1262–1271, 2001.
- D. Lam, D. Liu and T. Prescop. E-beam direct write (EBDW) as complementary lithography. In *Proceedings of SPIE*, volume 7823, page 78231C. 2010.
- M. Lam and A. Neureuther. A 3D substrate and buried defect simulator for EUV mask blanks. In *Proceedings of SPIE*, volume 5751, page 455. 2005.
- M. Land. *Evolutionary algorithms with local search for combinatorial optimization*. Ph.D. dissertation, University of California, San Diego, 1998.
- R. Landa Becerra and C. A. Coello Coello. Cultured differential evolution for constrained optimization. *Computer Methods in Applied Mechanics and Engineering*, 195(33-36):4303–4322, 2006.
- W. B. Langdon. Size Fair and Homologous Tree Genetic Programming Crossovers. *Genetic Programming and Evolvable Machines*, 1(1/2):95–119, 2000.
- P. Larranaga, C. Kuijpers, R. Murga, I. Inza and S. Dizdarevic. Genetic algorithms for the traveling salesman problem: A review of representations and operators. *Artificial Intelligence Review*, 13(2):129–170, 1999.
- M. Laumanns, L. Thiele and E. Zitzler. An efficient, adaptive parameter variation scheme for meta-heuristics based on the epsilon-constraint method. *European Journal of Operational Research*, 169(3):932–942, 2006.
- Y. LeCun, L. Bottou, G. Orr and K. Müller. Efficient BackProp. *Neural networks Tricks of the trade*, 1524(3):546–546, 1998.

- B. Lee, E. Hoshino, M. Takahashi, T. Yoneda, H. Yamanashi, H. Hoko, A. Chiba, M. Ito, T. Ogawa and S. Okazaki. Characteristics of Ru as a buffer layer or an etch stopper for EUVL mask patterning. *Microelectronic Engineering*, 61(1-4):233–239, 2002.
- M. A. Lee and H. Takagi. Dynamic Control of Genetic Algorithms using Fuzzy Logic Techniques. In *Proceedings of the 5th International Conference on Genetic Algorithms*, pages 76–83. Morgan Kaufmann, 1993.
- S. Lee, J. D. Byers, K. Jen, P. Zimmerman, B. Rice, N. J. Turro and C. G. Willson. An analysis of double exposure lithography options. In *Proceedings of SPIE*, volume 6924, pages 69242A–69242A–12. SPIE, 2008.
- J. Leonard, J. Carriere, J. Stack, R. Jones, M. D. Himel, J. Childers and K. Welch. An improved process for manufacturing diffractive optical elements (DOEs) for off-axis illumination systems. *Proceedings of SPIE*, 6924(2008):69242O–69242O–11, 2008.
- M. D. Levenson, J. S. Petersen, D. J. Gerold and C. A. Mack. Phase phirst! An improved strong-PSM paradigm. In B. J. Grenon and G. T. Dao (editors), *20th Annual BACUS Symposium on Photomask Technology*, volume 4186, pages 395–404. SPIE, Monterey, CA, USA, 2001.
- M. D. Levenson, N. Viswanathan and R. Simpson. Improving resolution in photolithography with a phase-shifting mask. *IEEE Transactions on Electron Devices*, 29(12):1828–1836, 1982.
- H. J. Levinson. *Principles of Lithography*. SPIE, Bellingham, WA, 2001.
- H. J. Levinson. *Principles of Lithography*. SPIE, Bellingham, WA, 2nd edition, 2005.
- J. Lewis, W. E. Hart and G. Ritchie. A comparison of dominance mechanisms and simple mutation on non-stationary problems. In *Parallel Problem Solving from Nature—PPSN V*, page 139. Springer, 1998.
- J. Li and Y. Lu. An Efficient Real-Coded Genetic Algorithm for Numerical Optimization Problems. In *Third International Conference on Natural Computation*, pages 760–764. IEEE Computer Society, 2007.
- L. Li. Use of Fourier series in the analysis of discontinuous periodic structures. *Journal of the Optical Society of America A*, 13(9):1870–1876, 1996.
- T. Li. Simulation of the Postexposure Bake Process of Chemically Amplified Resists by Reaction–Diffusion Equations. *Journal of Computational Physics*, 173(1):348–363, 2001.
- Y. Li, S.-M. Yu and Y.-L. Li. Intelligent optical proximity correction using genetic algorithm with model- and rule-based approaches. *Computational Materials Science*, 45(1):65–76, 2008.
- Y. Lian and X. Zhou. Fast and accurate computation of partially coherent imaging by stacked pupil shift operator. In *Proceedings of SPIE*, pages 74883G–74883G–10. 2009.
- T. Liang, E. Frendberg, B. Lieberman and A. Stivers. Advanced photolithographic mask repair using electron beams. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 23(6):3101, 2005.
- T. Liang and E. Gullikson. Multilayer Defect Compensation to Enable Quality Masks for EUVL Production. 2008.
- K. Lieberman, H. Terkel, M. Rudman, A. Ignatov and A. Lewis. High-resolution deep-UV laser mask repair based on near-field optical technology. In *Proceedings of SPIE*, volume 2793, page 481. SPIE, 1996.

- L. W. Liebmann. Layout impact of resolution enhancement techniques. In *Proceedings of the 2003 international symposium on Physical design - ISPD '03*, 3, page 110. ACM Press, New York, New York, USA, 2003.
- L. W. Liebmann, J. Bruce, W. Chu, M. Cross, I. Graur, J. Krueger, W. Leipold, S. M. Mansfield, A. McGuire and D. Sundling. Optimizing style options for subresolution assist features. In *Proceedings of SPIE*, volume 4346, page 141. SPIE, 2001.
- G. E. Liepins and M. D. Vose. Representational issues in genetic optimization. *Journal of Experimental & Theoretical Artificial Intelligence*, 2(2):101–115, 1990.
- B. J. Lin. The attenuated phase-shifting mask. *Solid State Technology*, 35(1):43–47, 1992a.
- B. J. Lin. The optimum numerical aperture for attenuated phase-shifting masks. *Microelectronic Engineering*, 17(1-4):79–85, 1992b.
- B. J. Lin. Deep uv lithography. *Journal of Vacuum Science and Technology*, 12(6):1317, 1975.
- B. J. Lin. The k_3 coefficient in nonparaxial λ/NA scaling equations for resolution, depth of focus, and immersion lithography. *Journal of Microlithography, Microfabrication, and Microsystems*, 1(1):7, 2002.
- B. J. Lin. Immersion lithography and its impact on semiconductor manufacturing. In B. W. Smith (editor), *Optical Microlithography XVII*, volume 5377, pages 46–67. SPIE, 2004.
- C.-M. Lin and W.-A. Loong. TiSixNy and TiSixOyNz as embedded materials for attenuated phase-shifting mask in 193 nm. *Microelectronic Engineering*, 46(1-4):93–96, 1999.
- J. Lin, J. Maul, N. Weber, C. Holfeld, M. Escher, M. Merkel, G. Schoenhense and U. Kleineberg. Inspection of EUVL mask blank defects and patterned masks using EUV photoemission electron microscopy. *Microelectronic Engineering*, 85(5-6):922–924, 2008.
- A. Lindenmayer. Mathematical models for cellular interactions in development I. Filaments with one-sided inputs. *Journal of Theoretical Biology*, 18(3):280–299, 1968a.
- A. Lindenmayer. Mathematical models for cellular interactions in development II. Simple and branching filaments with two-sided inputs. *Journal of Theoretical Biology*, 18(3):300–315, 1968b.
- S. H. Ling and F. H. F. Leung. An Improved Genetic Algorithm with Average-bound Crossover and Wavelet Mutation Operations. *Soft Computing*, 11(1):7–31, 2006.
- J. Lis and A. Eiben. *A multi-sexual genetic algorithm for multiobjective optimization*. IEEE, 1996.
- J. Lis and M. Lis. Self-adapting parallel genetic algorithm with the dynamic mutation probability, crossover rate and population size. In *Proceedings of the 1st Polish National Conference on Evolutionary Computation*, pages 324–329. Oficyna Wydawnicza Politechniki Warszawskiej, 1996.
- B. Lischke, W. Bennecke, M. Brunner, K. H. Herrmann, A. Heuberger, E. Knappek, P. Schäffer and U. Schnakenberg. Multi-Beam Concepts for Nanometer Devices. *Japanese Journal of Applied Physics*, 28(Part 1, No. 10):2058–2064, 1989.
- J. Liu. On setting the control parameter of the differential evolution method. In *Proc. 8th int. conf. soft computing (MENDEL 2002)*. 2002.
- P. Liu, Y. Cao, L. Chen, G. Chen, M. Feng, J. Jiang, H.-y. Liu, S. Suh, S.-W. Lee and S. Lee. Fast and accurate 3D mask model for full-chip OPC and verification. *Proceedings of SPIE*, 6520:65200R–65200R–12, 2007.

- S. Liu, T. Fühner, F. Shao, A. Barenbaum, J. Jahn and A. Erdmann. Topography-aware BARC optimization for double patterning. In *Proceedings of SPIE*, volume 7640, pages 76403C–76403C–12. 2010.
- S. Liu, G. Roeder, G. Aygun, K. Motzek, P. Evanschitzky and A. Erdmann. Simulation of 3D inclined/rotated UV lithography and its application to microneedles. *Optik - International Journal for Light and Electron Optics*, 123(10):928–931, 2012.
- Y. Liu, D. S. Abrams, L. Pang and A. Moore. Inverse lithography technology principles in practice: unintuitive patterns. In *Proceedings of SPIE*, volume 5992 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pages 886–893. 2005.
- Y. Liu and A. Zakhor. Binary and phase shifting mask design for optical lithography. *IEEE Transactions on Semiconductor Manufacturing*, 5(2):138–152, 1992.
- F. G. Lobo and C. F. Lima. Revisiting evolutionary algorithms with on-the-fly population size adjustment. In *Proceedings of the 8th annual conference on Genetic and Evolutionary Computation*, page 1241. ACM Press, New York, New York, USA, 2006.
- F. G. Lobo, C. F. Lima, Z. Michalewicz and K. A. De Jong. *Parameter Setting in Evolutionary Algorithms*, volume 54 of *Studies in Computational Intelligence*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2007.
- W.-A. Loong and C.-M. Lin. AlSixNy as an embedded layer for attenuated phase-shifting mask in 193 nm and the utilization of a chemically amplified negative resist NEB-22 for maskmaking. *Microelectronic Engineering*, 53(1-4):133–136, 2000.
- H. R. Lourenço, O. C. Martin and T. Stützle. Iterated Local Search: Framework and Applications. In M. Gendreau and J.-Y. Potvin (editors), *Handbook of Metaheuristics*, volume 146 of *International Series in Operations Research & Management Science*, chapter 12, pages 363–397. Springer, second edition, 2010.
- K. D. Lucas, H. Tanabe and A. J. Strojwas. Efficient and rigorous three-dimensional model for optical lithography simulation. *Journal of the Optical Society of America A*, 13(11):2187, 1996.
- K. D. Lucas, J. Word, G. Vandenbergh, S. Verhaegen and R. Jonckheere. Model-based OPC for first-generation 193-nm lithography. In *Proceedings of SPIE*, volume 4346, page 119. SPIE, 2001.
- C. Lucasius and G. Kateman. Application of genetic algorithms in chemometrics. In *Proceedings of the 3rd International Conference on Genetic Algorithms*, page 176. Morgan Kaufmann Publishers Inc., 1989.
- X. Ma and G. Arce. Binary mask optimization for inverse lithography with partially coherent illumination. *Journal of the Optical Society of America A*, 25(12):2960, 2008a.
- X. Ma and G. R. Arce. Generalized inverse lithography methods for phase-shifting mask design. *Opt. Express*, 15:15066–15079, 2007.
- X. Ma and G. R. Arce. PSM design for inverse lithography with partially coherent illumination. *Optics Express*, 16(24):20126, 2008b.
- D. Mac Donald, M. Nagler, C. Van Paski and T. Whitney. 160 Mpx/sec laser pattern generator for mask and reticle production. *Optical Microlithography III: Technology for the next decade, SATE*, 470:212–220, 1984.
- C. A. Mack. Development of Positive Photoresists. *Journal of The Electrochemical Society*, 134(1):148, 1987.

- C. A. Mack. Fundamental issues in phase-shifting mask technology. In *KTI Microlithography Seminar: Interface '91*, pages 23–25. 1991.
- C. A. Mack. New Kinetic Model for Resist Dissolution. *Journal of The Electrochemical Society*, 139(4):L35, 1992.
- C. A. Mack. Enhanced lumped parameter model for photolithography. *Proceedings of SPIE*, 2197:501–510, 1994.
- C. A. Mack. *Modeling solvent effects in optical lithography*. Ph.D. dissertation, University of Texas at Austin, 1998.
- C. A. Mack. Milestones in Optical Lithography Tool Suppliers. 2005.
- C. A. Mack. *Fundamental principles of optical lithography: the science of microfabrication*. Wiley-Interscience, 2007.
- C. A. Mack and G. Arthur. Notch Model for Photoresist Dissolution. *Electrochemical and Solid-State Letters*, 1(2):86, 1999.
- M. Maenhoudt, R. Gronheid, N. Stepanenko, T. Matsuda and D. Vangoidsenhoven. Alternative process schemes for double patterning that eliminate the intermediate etch step. In *Proceedings of SPIE*, volume 6924, page 69240P. 2008.
- M. Maenhoudt, J. Versluijs, H. Struyf, J. Van Olmen and M. Van Hove. Double patterning scheme for sub-0.25 μm single damascene structures at $\text{NA}=0.75$, $\lambda=193\text{nm}$. In *Proceedings of SPIE*, volume 5754, pages 1508–1518. SPIE, 2004.
- S. W. Mahfoud. Crowding and preselection revisited. Technical report, 1992.
- S. W. Mahfoud. *Niching Methods for Genetic Algorithms*. Ph.D. dissertation, University of Illinois at Urbana-Champaign, 1995.
- S. W. Mahfoud and D. E. Goldberg. Parallel recombinative simulated annealing: A genetic algorithm. *Parallel Computing*, 21(1):1–28, 1995.
- G. Y. Mak, E. Y. Lam and A. K.-K. Wong. Placement sensitivity to aberration in optical imaging. In *2003 IEEE Conference on Electron Devices and Solid-State Circuits (IEEE Cat. No.03TH8668)*, pages 475–478. IEEE, 2004.
- G. Y. Mak, A. K.-K. Wong and E. Y. Lam. Alternating phase-shifting mask design for low aberration sensitivity. *Journal of Microlithography, Microfabrication, and Microsystems*, 4(1):013008, 2005.
- R. A. Mäkinen, J. Periaux and J. Toivanen. Multidisciplinary shape optimization in aerodynamics and electromagnetics using genetic algorithms. *International Journal for Numerical Methods in Fluids*, 30(2):149–159, 1999.
- B. Manderick, M. de Weger and P. Spiessens. The genetic algorithm and the structure of the fitness landscape. In *Proceedings of the fourth international conference on genetic algorithms*, pages 143–150. San Mateo, CA: Morgan Kaufman, 1991.
- P. Mangat, S. Hector, S. Rose, G. Cardinale, E. Tejnill and A. Stivers. EUV mask fabrication with Cr absorber. In *Proceedings of SPIE*, volume 3997, page 76. SPIE, 2000.
- S. M. Mansfield, L. W. Liebmann, A. Molless and A. K.-K. Wong. Lithographic comparison of assist feature design strategies. In *Proceedings of SPIE*, volume 4000, page 63. SPIE, 2000.
- M. Mansuripur. Distribution of light at and near the focus of high-numerical-aperture objectives. *Journal of the Optical Society of America A*, 3(12):2086–2093, 1986.

- N. Maratos. *Exact penalty function algorithms for finite dimensional and control optimization problems*. Ph.D. dissertation, University of London, 1978.
- S. Marglin. *Public investment criteria*. The MIT Press, 1967.
- R. Martí, J. M. Moreno-Vega and A. Duarte. Advanced Multi-start Methods. In M. Gendreau and J.-Y. Potvin (editors), *Handbook of Metaheuristics*, volume 146 of *International Series in Operations Research & Management Science*, chapter 9, pages 265–281. Springer, second edition, 2010.
- M. Martínez, S. García-Nieto, J. Sanchis and X. Blasco. Genetic algorithms optimization for normalized normal constraint method under Pareto construction. *Advances in Engineering Software*, 40(4):260–267, 2009.
- L. Mashev and S. Tonchev. Formation of holographic diffraction gratings in photoresist. *Applied Physics A Solids and Surfaces*, 26(3):143–149, 1981.
- MathWorks. Optimization Toolbox for Use with MATLAB: User’s Guide, Version 3. 2004.
- D. Matiut, A. Erdmann, B. Tollkühn and A. Semmler. New models for the simulation of post-exposure bake of chemically amplified resists. In *Proceedings of SPIE*, pages 1132–1142. SPIE, 2003.
- K. Matsumoto, E. Costner, I. Nishimura, M. Ueda and C. G. Willson. High-index resist for 193-nm immersion lithography. In C. L. Henderson (editor), *Advances in Resist Materials and Processing Technology XXV*, volume 6923, page 692305. SPIE, 2008.
- T. Matsuo, K. Ohkubo, T. Haraguchi and K. Ueyama. Zr-based films for attenuated phase-shift mask. In *Proceedings of SPIE*, volume 3096, pages 354–361. SPIE, 1997.
- T. Matsuyama. The lithographic lens: its history and evolution. In *Proceedings of SPIE*, volume 6154, pages 615403–615403–14. SPIE, 2006.
- W. Maurer. Mask specifications for 193-nm lithography. In *Proceedings of SPIE*, volume 2884, pages 562–571. SPIE, 1996.
- M. McCord and R. Pease. Lithography with the scanning tunneling microscope. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 4(1):86–88, 1986.
- W. S. McCulloch and W. Pitts. A logical calculus of the ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics*, 5(4):115–133, 1943.
- M. McGeoch and C. Pike. Star pinch scalable EUV source. In *Proceedings of SPIE*, volume 5037, page 141. 2003.
- G. McIntyre, D. Sanders, R. Sooriyakumaran, H. Truong and R. Allen. The limitations of high-index resists for 193-nm hyper-NA lithography. In *Proceedings of SPIE*, volume 6923, pages 692304–692304–12. SPIE, 2008.
- K. I. M. McKinnon. Convergence of the Nelder–Mead Simplex Method to a Nonstationary Point. *SIAM Journal on Optimization*, 9(1):148, 1998.
- B. Meliorisz, S. Partel, T. Schnattinger, T. Fühner, A. Erdmann and P. Hudek. Investigation of high-resolution contact printing. *Microelectronic Engineering*, 85(5-6):744–748, 2008.
- J. Melngailis. Focused ion beam lithography. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 80-81:1271–1280, 1993.
- O. J. Mengshoel and D. E. Goldberg. Probabilistic crowding: Deterministic crowding with probabilistic replacement. In *Proceedings of the Genetic and*, volume I, pages 409–416. Morgan Kaufmann Publishers, San Francisco, CA, Orlando, FL, 1999.

- O. J. Mengshoel and D. E. Goldberg. The crowding approach to niching in genetic algorithms. *Evolutionary computation*, 16(3):315–54, 2008.
- P. Merz. A genetic local search approach to the quadratic assignment problem. In *Proceedings of the 7th International Conference on Genetic Algorithms*, pages 465–472. Morgan Kaufmann, 1997.
- A. Messac, A. Ismail-Yahaya and C. A. Mattson. The normalized normal constraint method for generating the Pareto frontier. *Structural and Multidisciplinary Optimization*, 25(2):86–98, 2003.
- Z. Michalewicz. *Genetic algorithms + data structures = evolution programs*. Springer-Verlag, third edition, 1996.
- Z. Michalewicz and N. Attia. Evolutionary optimization of constrained problems. In *Proceedings of the 3rd*, pages 98–108. 1994.
- Z. Michalewicz and C. Z. Janikow. Genetic algorithms for numerical optimization. *Statistics and Computing*, 1(2):75–91, 1991.
- Z. Michalewicz and M. Schoenauer. Evolutionary Algorithms for Constrained Parameter Optimization Problems. *Evolutionary Computation*, 4(1):1–32, 1996.
- Z. Michalewicz, G. A. Vignaux and M. Hobbs. A Nonstandard Genetic Algorithm for the Nonlinear Transportation Problem. *INFORMS Journal on Computing*, 3(4):307–316, 1991.
- A. Michel. *Abformung von Mikrostrukturen auf prozessierten Wafern*. Ph.D. dissertation, University Karlsruhe, Karlsruhe, 1992.
- K. Miettinen. *Nonlinear multiobjective optimization*, volume 12. Springer, Boston, Massachusetts, 1999.
- B. L. Miller and M. Shaw. Genetic algorithms with dynamic niche sharing for multimodal function optimization. In *Proceedings of IEEE International Conference on Evolutionary Computation*, pages 786–791. IEEE, 1996.
- S. C. Minne, J. D. Adams, G. Yaralioglu, S. R. Manalis, A. Atalar and C. F. Quate. Centimeter scale atomic force microscope imaging and lithography. *Applied Physics Letters*, 73(12):1742, 1998.
- M. Minsky and S. Papert. *Perceptrons: An Introduction to Computational Geometry*. MIT Press, Boston, MA, 1969.
- P. B. Mirkarimi, D. G. Stearns, S. L. Baker, J. W. Elmer, D. W. Sweeney and E. M. Gullikson. Method for repairing Mo/Si multilayer thin film phase defects in reticles for extreme ultraviolet lithography. *Journal of Applied Physics*, 91(1):81, 2002.
- M. G. Moharam, D. A. Pomett, E. B. Grann and T. K. Gaylord. Stable implementation of the rigorous coupled-wave analysis for surface-relief gratings: enhanced transmittance matrix approach. *Journal of the Optical Society of America A*, 12(5):1077, 1995.
- P. Monk. *Finite element methods for Maxwell's equations*. Oxford University Press, 2003.
- J. Moon, S. Yang, D. Pine and W. Chang. Multiple-exposure holographic lithography with phase shift. *Applied Physics Letters*, 85(18):4184–4186, 2004.
- G. E. Moore. Cramming more components onto integrated circuits. *Electronics*, 38(8):33–35, 1965.
- J. J. Moré and D. C. Sorensen. Computing a Trust Region Step. *SIAM Journal on Scientific and Statistical Computing*, 4(3):553–572, 1983.
- J. Morse. Reducing the size of the nondominated set: Pruning by clustering. *Computers & Operations Research*, 7(1-2):55–66, 1980.

- P. Moscato. On evolution, search, optimization, genetic algorithms and martial arts: Towards memetic algorithms. Technical Report C3P 826, California Institute of Technology, Pasadena, CA, 1989.
- S. Mostaghim and J. Teich. Quad-trees: A Data Structure for Storing Pareto Sets in Multiobjective Evolutionary Algorithms with Elitism. *Evolutionary Multiobjective Optimization*, pages 81–104, 2005.
- T. Motoki. Calculating the expected loss of diversity of selection schemes. *Evolutionary computation*, 10(4):397–422, 2002.
- H. Mühlenbein. Parallel Genetic Algorithms Population Genetics and Combinatorial Optimization. pages 416–421. Morgan Kaufmann Publishers Inc., 1989.
- H. Mühlenbein. Parallel genetic algorithms, population genetics and combinatorial optimization. In J. Becker, I. Eisele and F. Mündemann (editors), *Parallelism, Learning, Evolution*, volume 565 of *Lecture Notes in Computer Science*, pages 398–406. Springer Berlin / Heidelberg, 1991.
- H. Mühlenbein. How genetic algorithms really work: I. mutation and hillclimbing. *Parallel problem solving from nature*, 2:15–25, 1992.
- H. Mühlenbein, M. Gorges-Schleuter and O. Krämer. Evolution algorithms in combinatorial optimization. *Parallel Computing*, 7(1):65–85, 1988.
- H. Mühlenbein and D. Schlierkamp-Voosen. Predictive Models for the Breeder Genetic Algorithm I. Continuous Parameter Optimization. *Evolutionary Computation*, 1(1):25–49, 1993.
- H. Mühlenbein, M. Schomisch and J. Born. The parallel genetic algorithm as function optimizer. *Parallel Computing*, 17(6-7):619–632, 1991.
- M. Mukherjee, S. M. Mansfield, L. W. Liebmann, A. Lvov, E. Papadapoulou, M. Lavin and Z. Zhao. The problem of optimal placement of sub-resolution assist features (SRAF). In *Proceedings of SPIE*, volume 5754, page 1417. SPIE, 2005.
- M. Mulder, A. Engelen, O. Noordman, G. Streutker, B. van Driehuisen, C. van Nuenen, W. Endendijk, J. Verbeeck, W. Bouman, A. Bouma and Others. Performance of FlexRay: a fully programmable illumination system for generation of freeform sources on high NA immersion systems. In *Proceedings of SPIE*, volume 7640, page 76401P. 2010.
- T. Murata and H. Ishibuchi. MOGA: multi-objective genetic algorithms. In *Proceedings of the IEEE International Conference on Evolutionary Computation*, volume 1, page 289. IEEE, 1995.
- T. Mülders, W. Henke, K. Elian, C. Nölscher and M. Sebald. New stochastic post-exposure bake simulation method. *Journal of Microlithography, Microfabrication, and Microsystems*, 4(4):043010, 2005.
- S. Nahar and S. Sahni. Fast algorithm for polygon decomposition. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 7(4):473–483, 1988.
- K. H. Nakagawa, J. Chen, R. J. Socha, T. Laidig, K. Wampler, D. Van Den Broeke, M. Dusa and R. Caldwell. Halftone biasing OPC technology: an approach for achieving fine bias control on raster-scan systems. In *Proceedings of SPIE*, volume 3748, page 315. SPIE, 1999.
- Y. Nakajima, T. Sato, R. Inanami, T. Nakasugi and T. Higashiki. Aberration budget in extreme ultraviolet lithography. In *Proceedings of SPIE*, volume 6921, page 69211A. 2008.
- H. Nakamura, T. Shibata, K. Rikimaru, S. Ito, S. Tanaka and S. Inoue. Ion implantation as insoluble treatment for resist-stacking process. *Journal of Micro/Nanolithography, MEMS and MOEMS*, 9(1):013020, 2010.

- R. Nakano and T. Yamada. Conventional genetic algorithm for job shop problems. In *Proceedings of the Fourth International Conference on Genetic Algorithms*, pages 474–479. San Mateo, 1991.
- T. Namiki. A new FDTD algorithm based on alternating-direction implicit method. *IEEE Transactions on Microwave Theory and Techniques*, 47(10):2003–2007, 1999.
- K. M. Nashold and B. E. A. Saleh. Image construction through diffraction-limited high-contrast imaging systems: an iterative approach. *Journal of the Optical Society of America A*, 2(5):635, 1985.
- L. Nazareth and P. Tseng. Gilding the Lily: A Variant of the Nelder-Mead Algorithm Based on Golden-Section Search. *Computational Optimization and Applications*, 22:133–144, 2002.
- I. Nedelcu. *Interface structure and interdiffusion in Mo/Si multilayers*. Ph.D. dissertation, University of Twente, 2007.
- J. A. Nelder and R. Mead. A simplex method for function minimization. *The computer journal*, 7(4):308, 1965.
- A. Neureuther. Microlithography with soft x-rays. *Plenum Press, Synchrotron Radiation Research*, 1:223–275, 1980.
- K. P. Ng and K. C. Wong. A New Diploid Scheme and Dominance Change Mechanism for Non-Stationary Function Optimization. In *Proceedings of the 6th International Conference on Genetic Algorithms*, pages 159–166. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1995.
- G. Niimi, Y. Ueno, K. Nishigori, T. Aota, H. Yashiro and T. Tomie. Experimental evaluation of a stopping power of high-energy ions from a laser-produced plasma by a magnetic field. In *Proceedings of SPIE*, volume 5037, page 370. SPIE, 2003.
- N. I. Nikolaev and A. Erdmann. Rigorous simulation of alignment for microlithography. *Journal of Microlithography, Microfabrication, and Microsystems*, 2(3):220, 2003.
- J. L. Nistler, G. P. Hughes, A. J. Muray and J. N. Wiley. Issues associated with the commercialization of phase-shift masks. In K. C. McGinnis (editor), *11th Annual BACUS Symposium on Photomask Technology*, volume 1604, pages 236–264. SPIE, 1992.
- A. Nitayama, T. Sato, K. Hashimoto, F. Shigemitsu and M. Nakase. New phase shifting mask with self-aligned phase shifters for a quarter micron photolithography. In *Electron Devices Meeting, 1989. IEDM'89. Technical Digest., International*, pages 57–60. IEEE, 1989.
- A. E. Nix and M. D. Vose. Modeling genetic algorithms with Markov chains. *Annals of Mathematics and Artificial Intelligence*, 5(1):79–88, 1992.
- J. Nocedal. Updating quasi-Newton matrices with limited storage. *Mathematics of computation*, 35(151):773–782, 1980.
- J. Nocedal and S. J. Wright. *Numerical optimization*. Springer, New York, 1999.
- J. Nocedal and S. J. Wright. *Numerical optimization*. Springer, New York, second edition, 2006.
- T. Nomura and T. Miyoshi. Numerical coding and unfair average crossover in GA for fuzzy rule extraction in dynamic environments. In *Fuzzy Logic, Neural Networks, and Evolutionary Computation*, volume 1152, pages 55–72. Springer, 1996.
- P. Nordin. A compiling genetic programming system that directly manipulates the machine code. *Advances in genetic programming*, 1:311–331, 1994.
- A. B. J. Novikoff. On convergence proofs on perceptrons. In *Proceedings of the Symposium on the Mathematical Theory of Automata*, volume 12, pages 615–622. Polytechnic Institute of Brooklyn, Polytechnic Institute of Brooklyn, 1962.

- A. Nugrowati. *Vectorial diffraction of extreme ultraviolet light and ultrashort light pulses*. Ph.D. dissertation, TU Delft, 2008.
- D. Nyysönen and C. Kirk. Optical microscope imaging of lines patterned in thick layers with variable edge geometry: theory. *Journal of the Optical Society of America A*, 5(8):1270–1280, 1988.
- M. K. Oberthaler and T. Pfau. One-, two- and three-dimensional nanostructures with atom lithography. *Journal of Physics: Condensed Matter*, 15(6):R233–R255, 2003.
- C. K. Oei, D. E. Goldberg and S. J. Chang. Tournament selection, niching, and the preservation of diversity. Technical report, University of Illinois, 1991.
- Y.-H. Oh, J.-C. Lee and S. Lim. Resolution enhancement through optical proximity correction and stepper parameter optimization for 0.12- μm mask pattern. In *Proceedings of SPIE*, volume 3679, pages 607–613. Spie, 1999.
- H. Ohtsuka, T. Onodera, K. Kuwahara and T. Taguchi. Conjugate twin-shifter masks with multiple focal planes. In J. D. Cuthbert (editor), *Optical/Laser Microlithography V*, volume 1674, pages 53–64. SPIE, San Jose, CA, USA, 1992.
- U. Okoroanyanwu, J. Kye, N. Yamamoto and K. Cummings. Defectivity in water immersion lithography. *Microlithography World*, 14(4):4–7, 2005.
- W. Oldham. The use of contrast enhancement layers to improve the effective contrast of positive photoresist. *IEEE Transactions on Electron Devices*, 34(2):247–251, 1987.
- I. M. Oliver, D. J. Smith and J. R. C. Holland. A study of permutation crossover operators on the traveling salesman problem. In *Proceedings of the Second International Conference on Genetic Algorithms and Their Application*, pages 224–230. L. Erlbaum Associates Inc., Hillsdale, NJ, USA, 1987.
- M. Oltean. Evolving evolutionary algorithms using linear genetic programming. *Evolutionary Computation*, 13(3):387–410, 2005.
- Y.-S. Ong and A. Keane. Meta-Lamarckian Learning in Memetic Algorithms. *IEEE Transactions on Evolutionary Computation*, 8(2):99–110, 2004.
- Y.-S. Ong, M. Lim and N. Zhu. Classification of adaptive memetic algorithms: a comparative study. *Systems, Man, and*, 36(1):141–152, 2006.
- I. Ono, H. Satoh and S. Kobayashi. A real-coded genetic algorithm for function optimization using the unimodal normal distribution crossover. *Journal of Japanese Society for Artificial Intelligence*, 14:1146–1155, 1999.
- C. J. Ortiz, P. Pichler, T. Fühner, F. Cristiano, B. Colombeau, N. E. B. Cowern and A. Claverie. A physically based model for the spatial and temporal evolution of self-interstitial agglomerates in ion-implanted silicon. *Journal of Applied Physics*, 96(9):4866–4877, 2004.
- D. Orvosh and L. D. Davis. Using a genetic algorithm to optimize problems with feasibility constraints. In *Proceedings of the First IEEE Conference on Evolutionary Computation. IEEE World Congress on Computational Intelligence*, pages 548–553. IEEE, 1994.
- S. Osher and J. A. Sethian. Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulations. *Journal of Computational Physics*, 79(1):12–49, 1988.
- A. Osyczka and S. Kundu. A new method to solve generalized multicriteria optimization problems using the simple genetic algorithm. *Structural and Multidisciplinary Optimization*, 10(2):94–99, 1995.

- A. Otaka, Y. Kawai and Y. Sakakibara. Effect of lens aberration on hole pattern fabrication using halftone phase-shifting masks. In *Proceedings of SPIE*, volume 3096, pages 298–307. SPIE, 1997.
- K. Otaki. Asymmetric Properties of the Aerial Image in Extreme Ultraviolet Lithography. *Japanese Journal of Applied Physics*, 39(Part 1, No. 12B):6819–6826, 2000.
- T. Otsuka, D. Kilbane, J. White, T. Higashiguchi, N. Yugami, T. Yatagai, W. Jiang, A. Endo, P. Dunne and G. O’Sullivan. Rare-earth plasma extreme ultraviolet sources at 6.5–6.7 nm. *Applied Physics Letters*, 97(11):111503, 2010.
- O. Otto, J. Garofalo, K. Low, C. Yuan, R. Henderson, C. Pierrat, R. Kostelak, S. Vaidya and P. Vasudev. Automated optical proximity correction: a rules-based approach. In *Proceedings of SPIE*, volume 2197, page 278. SPIE, 1994.
- N. A. O’Connor, A. J. Berro, J. R. Lancaster, X. Gu, S. Jockusch, T. Nagai, T. Ogata, S. Lee, P. Zimmerman, C. G. Willson and N. J. Turro. Toward the Design of a Sequential Two Photon Photoacid Generator for Double Exposure Photolithography. *Chemistry of Materials*, 20(24):7374–7376, 2008.
- L. Pang, G. Dai, T. Cecil, T. Dam, Y. Cui, P. Hu, D. Chen, K.-H. Baik and D. Peng. Validation of inverse lithography technology (ILT) and its adaptive SRAF at advanced technology nodes. In *Proceedings of SPIE*, volume 6924, pages 69240T–69240T–12. Spie, 2008a.
- L. Pang, Y. Liu and D. S. Abrams. Inverse lithography technology (ILT): a natural solution for model-based SRAF at 45nm and 32nm. In *Proceedings of SPIE*, volume 6607, page 660739. SPIE, 2007.
- L. Pang, D. Peng, P. Hu, D. Chen, T. Cecil, L. He, G. Xiao, V. Tolani, T. Dam, K.-H. Baik and B. Gleason. Optimization from design rules, source and mask, to full chip with a single computational lithography framework: level-set-methods-based inverse lithography technology (ILT). In *Proceedings of SPIE*, pages 76400O–76400O–21. 2010.
- L. Pang, G. Xiao, V. Tolani, P. Hu, T. Cecil, T. Dam, K.-h. Baik and B. Gleason. Considering MEEF in Inverse Lithography Technology (ILT) and Source Mask Optimization (SMO). In *Proceedings of SPIE*, volume 7122 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pages 71221W–71221W–14. Spie, 2008b.
- J. Pankert, K. Bergmann, J. Klein, W. Neff, O. Rosier, S. Seiwert, C. Smith, S. Probst, D. Vaudrevange, G. Siemons and Others. Physical properties of the HCT EUV source. In *Proceedings of SPIE*, volume 5037, page 112. 2003.
- M. Papadrakakis, N. Lagaros and Y. Tsompanakis. Optimization of large-scale 3D trusses using Evolution Strategies and Neural Networks. *Int J Space Structures*, 14:211–223, 1999.
- M. Parikh. Corrections to proximity effects in electron beam lithography. I. Theory. *Journal of Applied Physics*, 50(6):4371, 1979.
- J. Park, C. Park, S. Rhie, Y. Kim, M. Yoo, J. Kong, H. Kim and S. Yoo. An efficient rule-based OPC approach using a DRC tool for 0.18 μm ASIC. In *Proceedings IEEE 2000 First International Symposium on Quality Electronic Design (Cat. No. PR00525)*, pages 81–85. IEEE Comput. Soc, 2000.
- J. M. Parkinson and D. Hutchinson. An Investigation into the Efficiency of Variants on the Simplex Method. In F. A. Lootsma (editor), *Numerical Methods for Nonlinear Optimization*, pages 115–135. Academic Press, 1972.
- L. Parthier, G. Wehrhan, D. Keutel, T. Aichele and M. Ansorg. Development progress of high refractive LuAG for hyper NA immersion systems. *Fourth International Symposium on Immersion Lithography*, 2007.

- W. Partlo, I. Fomenkov, I. Oliver and D. Birs. Development of an EUV (13.5 nm) light source employing a dense plasma focus in lithium vapor. In *Proceedings of SPIE*, volume 3997, page 136. 2000.
- A. Pascoletti and P. Serafini. Scalarizing vector optimization problems. *Journal of Optimization Theory and Applications*, 42(4):499–524, 1984.
- Y. C. Pati, A. A. Ghazanfarian and R. F. Pease. Exploiting structure in fast aerial image computation for integrated circuit patterns. *IEEE Transactions on Semiconductor Manufacturing*, 10(1):62–74, 1997.
- Y. C. Pati and T. Kailath. Phase-shifting masks for microlithography: automated design and mask requirements. *Journal of the Optical Society of America A*, 11(9):2438, 1994.
- K. Patterson, Y. Trouiller, K. D. Lucas, J. Belledent, A. Borjon, Y. Rody, C. Couderc, F. Sundermann, J. Urbani and S. Baron. Improving model-based OPC performance for the 65-nm node through calibration set optimization. In *Proceedings of SPIE*, volume 5756, page 294. SPIE, 2005.
- S. Peng, R. H. French, W. Qiu, R. C. Wheland, M. Yang, M. F. Lemon and M. K. Crawford. Second generation fluids for 193 nm immersion lithography. In *Proceedings of SPIE*, volume 5754, pages 427–434. 2005.
- S. Peng and William J. Adair. High-density plasma dry etch for DUV attenuated phase-shifting masks. In *Proceedings of SPIE*, volume 3412, pages 228–232. SPIE, 1998.
- P. Petric, C. Bevis, A. Carroll, H. Percy, M. Zywno, K. Standiford, A. Brodie, N. Bareket and L. Grella. REBL: A novel approach to high speed maskless electron beam direct write lithography. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 27(1):161, 2009.
- A. Petrowski. A clearing procedure as a niching method for genetic algorithms. In *Proceedings of IEEE International Conference on Evolutionary Computation ICEC-96*, pages 798–803. IEEE, 1996.
- A. K. Pfau, E. W. Scheckler, D. M. Newmark and A. R. Neureuther. Gradient phase-shifter transitions fabricated by ion milling. In *Proceedings of SPIE*, volume 1674, pages 585–593. SPIE, 1992.
- H. C. Pfeiffer. Variable spot shaping for electron-beam lithography. *Journal of Vacuum Science and Technology*, 15(3):887, 1978.
- D. T. Pham, A. Ghanbarzadeh and E. Koc. The Bees Algorithm – A Novel Tool for Complex Optimization Problems. In *Proceedings of IPROMS*, pages 454–459. 2006.
- S. P. Phelps and M. Köksalan. An Interactive Evolutionary Metaheuristic for Multiobjective Combinatorial Optimization. *Management Science*, 49(12):1726–1738, 2003.
- C. Pierrat, A. K.-K. Wong and S. Vaidya. Phase-shifting mask topography effects on lithographic image quality. In *Electron Devices Meeting, 1992. Technical Digest., International*, pages 53–56. IEEE, 1992.
- S. Pierret and R. Van Den Braembussche. Turbomachinery Blade Design Using a Navier-Stokes Solver and Artificial Neural Network. *Journal of Turbomachinery*, 121:326–332, 1999.
- R. D. Piner. "Dip-Pen" Nanolithography. *Science*, 283(5402):661–663, 1999.
- T. Pistor, Y. Deng and A. Neureuther. Extreme ultraviolet mask defect simulation: Low-profile defects. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 18:2926, 2000a.
- T. Pistor, A. Neureuther and R. J. Socha. Modeling oblique incidence effects in photomasks. In *Proceedings of SPIE*, volume 4000, pages 228–237. 2000b.

- E. Platzgummer. Maskless lithography and nanopatterning with electron and ion multi-beam projection. In *Proceedings of SPIE*, volume 7637, page 763703. 2010.
- E. Platzgummer, H. Loeschner and G. Gross. Projection maskless patterning (PMLP) for the fabrication of leading-edge complex masks and nano-imprint templates. In *Proceedings of SPIE*, volume 6730, page 673033. SPIE, 2007.
- D. W. Pohl, W. Denk and M. Lanz. Optical stethoscopy: Image recording with resolution $\lambda/20$. *Applied Physics Letters*, 44(7):651, 1984.
- R. Poli. Tournament selection, iterated coupon-collection problem, and backward-chaining evolutionary algorithms. *Foundations of Genetic Algorithms*, pages 382–387, 2005.
- R. Poli and W. B. Langdon. A New Schema Theory for Genetic Programming with One-point Crossover and Point Mutation. In *Genetic Programming 1997: Proceedings of the Second Annual Conference*, pages 278–285. Morgan Kaufmann, 1997.
- R. Poli and W. B. Langdon. On the Search Properties of Different Crossover Operators in Genetic Programming. In *Genetic Programming 1998: Proceedings of the Third Annual Conference*, pages 293–301. Morgan Kaufmann, 1998.
- R. Poli, W. B. Langdon and N. McPhee. *A field guide to genetic programming*. Lulu Enterprises Uk Ltd, 2008.
- A. Poonawala and P. Milanfar. OPC and PSM design using inverse lithography: a nonlinear optimization approach. In *Proceedings of SPIE*, volume 6154, pages 61543H–61543H–14. SPIE, 2006.
- A. Poonawala and P. Milanfar. A pixel-based regularization approach to inverse lithography. *Micro-electronic Engineering*, 84(12):2837–2852, 2007a.
- A. Poonawala and P. Milanfar. Double-exposure mask synthesis using inverse lithography. *Journal of Micro/Nanolithography, MEMS and MOEMS*, 6(4):043001, 2007b.
- A. Poonawala and P. Milanfar. Mask design for optical microlithography—an inverse imaging problem. *IEEE transactions on image processing : a publication of the IEEE Signal Processing Society*, 16(3):774–88, 2007c.
- O. Popa and R. Jonckheere. Attenuated phase shift masks using MoSi as an opaque layer. In *Semiconductor Conference, 1995. CAS'95 Proceedings., 1995 International*, pages 171–174. IEEE, 1995.
- S. Popp. Entwicklung einer Umgebung zur fehlertoleranten Verteilung rechenintensiver Aufgaben in Python. Praktikumsbericht, Fraunhofer IISB, Erlangen, Germany, 2005.
- S. Popp. *Funktionsapproximation in memetischen Algorithmen mit Hilfe künstlicher neuronaler Netze*. Diplomarbeit, University of Applied Sciences Regensburg, 2007.
- M. Powell. A fast algorithm for nonlinearly constrained optimization calculations. *Numerical analysis*, 630:144–157, 1978a.
- M. Powell. The convergence of variable metric methods for nonlinearly constrained optimization calculations. *Nonlinear programming*, 3:27–63, 1978b.
- M. Powell and P. Toint. On the estimation of sparse Hessian matrices. *SIAM Journal on Numerical Analysis*, 1979.
- K. Price. Differential evolution: a fast and simple numerical optimizer. In *Proceedings of North American Fuzzy Information Processing*, pages 524–527. IEEE, 1996.

- M. Prouty and A. R. Neureuther. Optical imaging with phase shift masks. *SPIE milestone series*, 178:397–401, 2004.
- A. Prügel-Bennett and J. L. Shapiro. The dynamics of a Genetic Algorithm for simple random Ising systems. *Physica D: Nonlinear Phenomena*, 104(1):75–114, 1997.
- X. Qi and F. Palmieri. Theoretical analysis of evolutionary algorithms with an infinite population size in continuous space. Part I: Basic properties of selection and mutation. *IEEE transactions on neural networks: a publication of the IEEE Neural Networks Council*, 5(1):102–19, 1994a.
- X. Qi and F. Palmieri. Theoretical analysis of evolutionary algorithms with an infinite population size in continuous space. Part II: Analysis of the diversification role of crossover. *IEEE transactions on neural networks: a publication of the IEEE Neural Networks Council*, 5(1):120–9, 1994b.
- N. J. Radcliffe. Equivalence class analysis of genetic algorithms. *Complex Systems*, 5(2):183–205, 1991.
- N. J. Radcliffe. Non-Linear Genetic Representations. *Parallel Problem Solving from Nature 2*, pages 259–268, 1992.
- N. J. Radcliffe and P. Surry. Fundamental limitations on search algorithms: Evolutionary computing in perspective. *Computer Science Today*, pages 275–291, 1995.
- N. J. Radcliffe and P. D. Surry. Formal memetic algorithms. *Evolutionary computing*, 865:1–16, 1994.
- Z. Rahimi, A. Erdmann, P. Evanschitzky and C. Pflaum. Rigorous EMF simulation of absorber shape variations and their impact on lithographic processes. In *Proceedings of SPIE*, volume 7545, pages 75450C–75450C–10. 2010.
- J. N. Randall. High resolution ion beam lithography at large gaps using stencil masks. *Applied Physics Letters*, 42(5):457, 1983.
- S. Ranji Ranjithan, S. Kishan Chetan and H. Dakshina. Constraint Method-Based Evolutionary Algorithm (CMEA) for Multiobjective Optimization. In E. Zitzler, L. Thiele, K. Deb, C. Coello Coello and D. Corne (editors), *Evolutionary Multi-Criterion Optimization*, volume 1993 of *Lecture Notes in Computer Science*, pages 299–313. Springer Berlin / Heidelberg, 2001.
- M. Ratford, A. Tuson and H. Thompson. Applying Sexual Selection as a Mechanism for Obtaining Multiple Distinct Solutions. Technical report, The University of Edingburgh, 1997.
- A. Raub, A. Frauenglass, S. Brueck, W. Conley, R. Dammel, A. Romano, M. Sato and W. Hinsberg. Deep-UV immersion interferometric lithography. In *Proceedings of SPIE*, volume 5377, page 306. SPIE, 2004.
- L. Rayleigh. Investigations in optics, with special reference to the spectroscope, B 7. *SPIE milestone series*, 178:43–50, 2004.
- I. Rechenberg. *Evolutionsstrategie: Optimierung technischer Systeme nach Prinzipien der biologischen Evolution*, Fromman-Holzboog. Frommann-Holzboog Verlag, Stuttgart, 1973.
- C. R. Reeves. A genetic algorithm for flowshop sequencing. *Computers & Operations Research*, 22(1):5–13, 1995.
- C. Reynolds. Flocks, herds and schools: A distributed behavioral model. In *ACM SIGGRAPH Computer Graphics*, volume 21, pages 25–34. ACM, 1987.
- G. A. M. Reynolds, R. H. French, P. F. Carcia, C. C. Torardi, G. P. Hughes, D. J. Jones, M. F. Lemon, M. H. Reilly, L. Wilson and C. R. Miao. TiSi-nitride attenuating phase-shift photomask for 193-nm lithography. *Proceedings of SPIE*, 3546:514–523, 1998.

- J. Richardson, M. Palmer, G. E. Liepins and M. Hilliard. Some guidelines for genetic algorithms with penalty functions. In *Proceedings of the 3rd International Conference on Genetic Algorithms*, pages 191–197. Morgan Kaufmann Publishers Inc., 1989.
- M. Riedmiller and H. Braun. RPROP - A Fast Adaptive Learning Algorithm. In *Proceedings of the International Symposium on Computer and Information Science VII*. 1992.
- M. Rieger, J. Mayhew and S. Panchapakesan. Layout design methodologies for sub-wavelength manufacturing. In *Proceedings of the 38th Design Automation Conference (IEEE Cat. No.01CH37232)*, pages 85–88. ACM, 2001.
- RIT Nanolithography Research Labs. Optical Properties of Thin Films. 2012.
- S. A. Robertson, E. Pavelcheck, W. Hoppe and R. Wildfeuer. Improved notch model for resist dissolution in lithography simulation. In *Proceedings of SPIE*, volume 4345, pages 912–920. Spie, 2001.
- J. Robic, P. Schiavone, V. Rodillon and R. Payerne. Optimisation of EUV mask absorbing layers. *Microelectronic Engineering*, 61:257–263, 2002.
- R. Rojas. *Neural networks: a systematic introduction*. Springer, 1996.
- C. Romero, M. Tamiz and D. Jones. Goal programming, compromise programming and reference point method formulations: linkages and utility interpretations. *Journal of the Operational Research Society*, 49(9):986–991, 1998.
- E. Ronald. When selection meets seduction. In *Proceedings of the 6th International Conference on Genetic Algorithms*, pages 167–173. Morgan Kaufmann Publishers Inc., 1995.
- K. Ronse, R. Pforr, K.-H. Baik, R. Jonckheere and L. Van Den Hove. Attenuated phase shifting masks in combination with off-axis illumination: a way towards quarter micron DUV lithography for random logic applications. *Microelectronic Engineering*, 23(1-4):133–138, 1994.
- J. P. Rosca and D. H. Ballard. Hierarchical Self-Organization in Genetic Programming. In *Proceedings of the Eleventh International Conference on Machine Learning*. The MIT Press, 1994.
- R. S. Rosenberg. *Simulation of genetic populations with biochemical properties*. Ph.D. dissertation, University of Michigan, Ann Harbor, Michigan, 1967.
- F. Rosenblatt. The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, 65(6):386–408, 1958.
- A. E. Rosenbluth. Global optimization of the illumination distribution to maximize integrated process window. In *Proceedings of SPIE*, volume 6154, pages 61540H–61540H–12. SPIE, 2006.
- A. E. Rosenbluth, S. Bukofsky, C. Fonseca, M. Hibbs, K. Lai, A. Molless, R. Singh and A. K.-K. Wong. Optimum mask and source patterns to print a given shape. *Journal of Microlithography, Microfabrication, and Microsystems*, 1(1):13, 2002.
- A. E. Rosenbluth, D. Melville, K. Tian, K. Lai, N. Seong, D. Pfeiffer and M. Colburn. Global optimization of masks, including film stack design to restore TM contrast in high NA TCC's. In *Proceedings of SPIE*, volume 6520, pages 65200P–65200P–13. SPIE, 2007.
- H. H. Rosenbrock. An Automatic Method for Finding the Greatest or Least Value of a Function. *The Computer Journal*, 3(3):175–184, 1960.
- R. Rosenthal. Concepts, Theory, and Techniques PRINCIPLES OF MULTIOBJECTIVE OPTIMIZATION. *Decision Sciences*, 16(2):133–152, 1985.
- F. Rothlauf. *Representations for genetic and evolutionary algorithms*. Springer Verlag, 2006.

- J. E. Rowe and I. East. Deception and GA-hardness. In *Proceedings of the First Nordic Workshop on Genetic Algorithms and Their Applications*, pages 165–172. 1995.
- G. Rudolph. Convergence analysis of canonical genetic algorithms. *Neural Networks, IEEE Transactions on*, 5(1):96–101, 1994.
- G. Rudolph. Convergence of evolutionary algorithms in general search spaces. In *Proceedings of IEEE International Conference on Evolutionary Computation*, pages 50–54. IEEE, 1996.
- D. E. Rumelhart, G. E. Hinton and R. J. Williams. Learning internal representation by error propagation. In D. E. Rumelhart and J. L. McClelland (editors), *Parallel Distributed Processing Explorations in the Microstructure of Cognition*, volume 1, pages 318–362. MIT Press, Cambridge, MA, 1986.
- J. Ruoff, J. Neumann, E. Schmitt-Weaver, E. Setten, N. Masson, C. Proglar and B. Geh. Polarization-induced astigmatism caused by topographic masks. In *Mask and Lithography Conference (EMLC), 2008 24th European*, pages 1–15. VDE, 2008.
- C. Ryan. The degree of oneness. In *Proceedings of the First Online Workshop on Soft Computing (WSC1)*, pages 43–48. Citeseer, 1994.
- K. Sakajiri, A. Trichkov, Y. Granik, E. Hendrickx, G. Vandenbergh, M. Kempsell, G. Fenger, K. Boehm and T. Scheruebl. Application of pixel-based mask optimization technique for high transmission attenuated PSM. In *Proceedings of SPIE*, volume 7275, pages 72750X–72750X–11. SPIE, 2009.
- R. Sakamoto, T. Endo, B.-C. Ho, S. Kimura, T. Ishida, M. Kato, N. Fujitani, R. Onishi, Y. Hiroi and D. Maruyama. Bottom anti-reflective coatings (BARC) for LFLE double patterning process. In *Proceedings of SPIE*, volume 7639, pages 76391R–76391R–10. 2010.
- K. Salaita, Y. Wang, J. Fragala, R. Vega, C. Liu and C. Mirkin. Massively Parallel Dip-Pen Nanolithography with 55,000-Pen Two-Dimensional Arrays. *Angewandte Chemie*, 118(43):7378–7381, 2006.
- B. E. A. Saleh and K. M. Nashold. Image construction: optimum amplitude and phase masks in photolithography. *Appl. Opt.*, 24(10):1432–1437, 1985.
- B. E. A. Saleh and M. Rabbani. Simulation of partially coherent imagery in the space and frequency domains and by modal expansion. *Applied optics*, 21(15):2770–7, 1982.
- B. E. A. Saleh and S. I. Sayegh. Reduction of errors of microphotographic reproductions by optimal corrections of original masks. *Optical Engineering*, 20:781–787, 1981.
- J. Sánchez-Velazco. A Gendered Selection Strategies in Genetic Algorithms for Optimization. In *Proceedings of the UK Workshop on Computational Intelligence*, pages 217–223. University of Bristol, 2003.
- A. Sasaki, K. Nishihara, A. Sunahara, H. Furukawa, T. Nishikawa and F. Koike. Theoretical investigation of the spectrum and conversion efficiency of short wavelength extreme-ultraviolet light sources based on terbium plasmas. *Applied Physics Letters*, 97(23):231501, 2010.
- K. Sato, M. Itoh and T. Sato. Mask 3D effect on 45-nm imaging using attenuated PSM. In *Proceedings of SPIE*, volume 6520, page 65201J. 2007.
- S. I. Sayegh. *Image restoration and image design in nonlinear optical systems*. Ph.D. dissertation, University of Wisconsin, Madison, 1982.
- J. Schaffer and L. J. Eshelman. On crossover as an evolutionarily viable strategy. In *Proceedings of the Fourth International Conference on Genetic Algorithms*, pages 61–68. Morgan Kaufmann Publishers, 1991.

- J. D. Schaffer. Multiple objective optimization with vector evaluated genetic algorithms. In *Proceedings of the 1st International Conference on Genetic Algorithms*, pages 93–100. L. Erlbaum Associates Inc., 1985.
- J. D. Schaffer, R. Caruana, L. J. Eshelman and R. Das. A study of control parameters affecting online performance of genetic algorithms for function optimization. In *Proceedings of the third international conference on Genetic algorithms*, page 60. Morgan Kaufmann Publishers Inc., 1989.
- J. D. Schaffer and A. Morishima. An adaptive crossover distribution mechanism for genetic algorithms. In *Proceedings of the Second International Conference on Genetic Algorithms on Genetic algorithms and their application*, pages 36–40. L. Erlbaum Associates Inc., Hillsdale, NJ, USA, 1987.
- F. M. Schellenberg, O. Toublan, L. Capodieci and R. J. Socha. Adoption of OPC and the impact on design and layout. In *Proceedings of the 38th conference on Design automation - DAC '01*, pages 89–92. ACM Press, New York, New York, USA, 2001.
- R. Schenker, S. Bollepalli, B. Hu, K. K. Toh, V. Singh, K. Yung, W.-h. Cheng and Y. Borodovsky. Integration of pixelated phase masks for full-chip random logic layers. In *Proceedings of SPIE*, volume 6924, pages 69240I–69240I–11. SPIE, 2008.
- P. Schiavone, G. Granet and J. Robic. Rigorous electromagnetic simulation of EUV masks: influence of the absorber properties. *Microelectronic Engineering*, 57:497–503, 2001.
- K. Schittkowski and Y.-X. Yuan. *Sequential Quadratic Programming Methods*. John Wiley & Sons, Inc., 2011.
- D. Schlierkamp-Voosen and H. Mühlenbein. Adaptation of population sizes by competing subpopulations. In *Evolutionary Computation, 1996., Proceedings of IEEE International Conference on*, pages 330–335. IEEE, 1996.
- T. Schnattinger. *Mesoscopic simulation of photoresist processing in optical lithography*. Ph.D. dissertation, University Erlangen-Nuremberg, 2007.
- T. Schnier and X. Yao. Using multiple representations in evolutionary algorithms. In *Evolutionary Computation, 2000. Proceedings of the 2000 Congress on*, volume 1, pages 479–486. IEEE, 2000.
- C. Schomburg, B. Höfflinger, R. Springer and R. Wijnaendts-van Resandt. Economic production of submicron ASICs with laser beam direct write lithography. *Microelectronic Engineering*, 35(1-4):509–512, 1997.
- J. Schott. *Fault Tolerant Design Using Single and Multicriteria Genetic Algorithm Optimization*. Master thesis, Massachusetts Institute of Technology, 1995.
- N. Schraudolph and R. Belew. Dynamic parameter encoding for genetic algorithms. *Machine Learning*, 9(1):9–21, 1992.
- H.-P. Schwefel. *Numerische Optimierung von Computer Modellen mittels der Evolutionsstrategie*. Birkhäuser, Basel, 1977.
- H.-P. Schwefel. Collective intelligence in evolving systems. *Ecodynamics, Contributions to Theoretical Ecology*, pages 95–100, 1987.
- M. Sears, G. Fenger, J. Mailfert and B. Smith. Extending SMO into the lens pupil domain. In *Proceedings of SPIE*, volume 7973, page 79731B. 2011.
- D. Sedighizadeh and E. Masehian. Particle Swarm Optimization Methods, Taxonomy and Applications. *International Journal of Computer Theory and Engineering*, 1(5):1793–8201, 2009.
- S. Seifert. An Introduction to Pythmea. User’s manual, Fraunhofer IISB, Erlangen, Germany, 2006a.

- S. Seifert. *Multi-objective Optimization of Lithographic Process Conditions Using a Genetic Algorithm*. Master thesis, University Erlangen-Nuremberg, 2006b.
- S. Seifert. OptProb 0.1 – A generic interface for the specification of optimization problems. User's manual, Fraunhofer IISB, Erlangen, Germany, 2007.
- R. L. Seliger and W. P. Fleming. Focused ion beams in microfabrication. *Journal of Applied Physics*, 45(3):1416, 1974.
- D. Seniw. *A non-standard genetic algorithm for the traveling salesman problem*. Ph.D. dissertation, University of North Carolina at Charlotte, 1991.
- J. A. Sethian. A fast marching level set method for monotonically advancing fronts. *Proceedings of the National Academy of Sciences of the United States of America*, 93(4):1591, 1996.
- J. A. Sethian and D. Adalsteinsson. An overview of level set methods for etching, deposition, and lithography development. *IEEE Transactions on Semiconductor Manufacturing*, 10(1):167–184, 1997.
- H. Sewell, J. Mulken, P. Graeupner, D. McCafferty, L. Markoya, S. Donders, N. Samarakone and R. Duesing. Extending Immersion Lithography with High Index Materials Results of a feasibility study. In *Proceedings of SPIE*, volume 6520, page 55. 2007.
- C. Shaefer. The ARGOT strategy: Adaptive representation genetic optimizer technique. In *Proceedings of the Second International Conference on Genetic Algorithms*, pages 50–58. 1987.
- D. F. Shanno. Conditioning of quasi-Newton methods for function minimization. *Mathematics of Computation*, 24:647–656, 1970.
- F. Shao, G. Cooper, Z. Chen and A. Erdmann. Modeling of exploration of reversible contrast enhancement layers for double exposure lithography. In *Proceedings of SPIE*, volume 7640, page 76400J. 2010.
- F. Shao, P. Evanschitzky, T. Fühner and A. Erdmann. Rigorous diffraction simulations of topographic wafer stacks in double patterning. *Microelectronic Engineering*, 86(4-6):489–491, 2009.
- F. Shao, P. Evanschitzky, D. Reibold and A. Erdmann. Fast rigorous simulation of mask diffraction using the waveguide method with parallelized decomposition technique. In *Proceedings of SPIE*, volume 6792, pages 679206–679206–10. SPIE, 2008.
- J. L. Shapiro, A. Prügel-Bennett and M. Rattray. A Statistical Mechanical Formulation of the Dynamics of Genetic Algorithms. In *Evolutionary Computing, AISB Workshop*, pages 17–27. 1994.
- J. Shekel. Test functions for multimodal search techniques. In *Fifth Annual Princeton Conference on Information Science and Systems*, pages 354–359. 1971.
- Y. Shen, N. Wong and E. Y. Lam. Level-set-based inverse lithography for photomask synthesis. *Optics Express*, 17(26):23690, 2009.
- S. Sherif, B. E. A. Saleh and R. De Leone. Binary image synthesis using mixed linear integer programming. *IEEE transactions on image processing : a publication of the IEEE Signal Processing Society*, 4(9):1252–7, 1995.
- M. Shibuya. Original Plate To Be Projected For Use In Transmission Illumination. Japanese Patent Number 57-62052, 1982.
- H. Shields, S. Fornaca, M. Petach, M. Michaelian, R. McGregor, R. Moyer and R. Pierre. Xenon target performance characteristics for laser-produced plasma EUV sources. In *Proceedings of SPIE*, volume 4688, page 94. SPIE, 2002.

- H.-D. Shin, C. Y. Jeoung, T. G. Kim, S. Lee, I.-S. Park and J. Ahn. Effect of Attenuated Phase Shift Mask Structure on Extreme Ultraviolet Lithography. *Japanese Journal of Applied Physics*, 48(No. 6):06FA06, 2009.
- B. W. Silverman. *Density estimation for statistics and data analysis*. Chapman and Hall/CRC Press, 1986.
- J. P. Silverman. Challenges and progress in x-ray lithography. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 16(6):3137, 1998.
- M. Singh and J. Braat. Improved theoretical reflectivities of extreme-ultraviolet mirrors. In *Proc. SPIE*, volume 3997, pages 412–419. SPIE, 2000.
- V. Singh, B. Hu, K. K. Toh, S. Bollepalli, S. Wagner and Y. Borodovsky. Making a trillion pixels dance. In *Proceedings of SPIE*, volume 6924, pages 69240S–69240S–12. SPIE, 2008.
- S. D. Slonaker. Further pursuit of correlation between lens aberration content and imaging performance. In *Proceedings of SPIE*, volume 4346, pages 1394–1403. SPIE, 2001.
- A. E. Smith and D. Tate. Genetic optimization using a penalty function. In *Proceedings of the 5th International Conference on Genetic Algorithms*, pages 499–505. Morgan Kaufmann Publishers Inc., 1993.
- B. W. Smith, S. Butt, Z. Alam, S. Kurinec and R. Lane. Attenuated phase shift mask materials for 248 and 193 nm lithography. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 14(6):3719–3723, 1996.
- B. W. Smith. Understanding lens aberration and influences to lithographic imaging. *Proceedings of SPIE*, 4000:294–306, 2000.
- B. W. Smith. Mutual Optimization of Resolution Enhancement Techniques. *Journal of Microlithography, Microfabrication, and Microsystems*, 1:95, 2002.
- B. W. Smith, A. Bourov, H. Kang, F. Cropanese, Y. Fan, N. Lafferty and L. Zavyalova. Water immersion optical lithography at 193 nm. *Journal of Microlithography, Microfabrication, and Microsystems*, 3:44, 2004.
- B. W. Smith and L. Zavyalova. Illumination pupil filtering using modified quadrupole apertures. *Proceedings of SPIE*, 3334:384–394, 1998.
- J. E. Smith and T. C. Fogarty. Self adaptation of mutation rates in a steady state genetic algorithm. In *Evolutionary Computation, 1996., Proceedings of IEEE International Conference on*, pages 318–323. IEEE, 1996.
- J. E. Smith. Coevolving Memetic Algorithms: A Review and Progress Report. *IEEE Transactions on Systems, Man and Cybernetics, Part B (Cybernetics)*, 37(1):6–17, 2007.
- J. M. Smith. When learning guides evolution. *Nature*, 329(6142):761–2, 1987.
- R. E. Smith, S. Forrest and A. S. Perelson. Searching for Diverse, Cooperative Populations with Genetic Algorithms. *Evolutionary Computation*, 1(2):127–149, 1993.
- R. E. Smith and D. E. Goldberg. Diploidy and dominance in artificial genetic search. *Complex Systems*, 6(3):251–285, 1992.
- R. E. Smith and E. Smuda. Adaptively resizing populations: Algorithm, analysis, and first results. *Complex Systems*, 9, 1995.

- E. S. Snow and P. M. Campbell. Fabrication of Si nanostructures with an atomic force microscope. *Applied Physics Letters*, 64(15):1932, 1994.
- R. J. Socha, X. Shi and D. LeHoty. Simultaneous source mask optimization (SMO). In *Proceedings of SPIE*, volume 5853, pages 180–193. SPIE, 2005.
- R. J. Socha, D. Van Den Broeke, S. Hsu, J. Chen, T. Laidig, N. Corcoran, U. Hollerbach, K. Wampler, X. Shi and W. Conley. Contact hole reticle optimization by using interference mapping lithography (IML). In *Proceedings of SPIE*, volume 5377, page 222. SPIE, 2004.
- A. Sokolov and L. D. Whitley. Unbiased tournament selection. In *Proceedings of the 2005 Conference on Genetic and Evolutionary Computation*, page 1131. ACM Press, New York, New York, USA, 2005.
- H. H. Solak, C. David, J. Gobrecht, L. Wang and F. Cerrina. Four-wave EUV interference lithography. *Microelectronic Engineering*, 61(1-3):77–82, 2002.
- H. H. Solak, D. He, W. Li and F. Cerrina. Nanolithography using extreme ultraviolet lithography interferometry: 19 nm lines and spaces. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 17(6):3052, 1999.
- D. Spears and H. I. Smith. High-resolution pattern replication using soft X rays. *Electronics Letters*, 8(4):102, 1972.
- W. M. Spears. Crossover or mutation. *Foundations of genetic algorithms*, 2:221–237, 1993.
- W. M. Spears. Adapting crossover in evolutionary algorithms. *Evolutionary programming 4*, page 367, 1995.
- W. M. Spears. *The role of mutation and recombination in evolutionary algorithms*. Ph.D. dissertation, George Mason University, 1998.
- W. M. Spears and K. A. De Jong. On the virtues of parameterized uniform crossover. In *Proceedings of the Fourth International Conference on Genetic Algorithms*, pages 230–236. 1991.
- W. M. Spears and K. A. D. Jong. An analysis of multipoint crossover. In *Proc of the Foundation Of Genetic Algorithms*, AIC-90-014, pages 301–315. Navy Center for Applied Research in Artificial Intelligence, Morgan Kaufmann (San Mateo), 1991.
- P. Spellucci. A new technique for inconsistent QP problems in the SQP method. *Mathematical Methods of Operations Research*, 47(3):355–400, 1998a.
- P. Spellucci. An SQP method for general nonlinear programs using only equality constrained subproblems. *Mathematical programming*, 82(3):413–448, 1998b.
- W. Spendley, G. R. Hext and F. R. Himsworth. Sequential Application of Simplex Designs in Optimisation and Evolutionary Operation. *Technometrics*, 4(4):441–461, 1962.
- M. Srinivas and L. Patnaik. Adaptive probabilities of crossover and mutation in genetic algorithms. *IEEE Transactions on Systems, Man, and Cybernetics*, 24(4):656–667, 1994.
- N. Srinivas and K. Deb. Multiobjective Optimization Using Nondominated Sorting in Genetic Algorithms. *Evolutionary Computation*, 2(3):221–248, 1994.
- W. Srituravanich, L. Pan, Y. Wang, C. Sun, D. B. Bogy and X. Zhang. Flying plasmonic lens in the near field for high-speed nanolithography. *Nature nanotechnology*, 3(12):733–7, 2008.
- U. Stamm, I. Ahmad, I. Balogh, H. Birner, D. Bolshukhin, J. Bruderhmann, S. Enke, F. Flohrer, G. Kai, G. Hergenhan and Others. High-power EUV lithography sources based on gas discharges and laser-produced plasmas. In *Proceedings of SPIE*, volume 5037, page 119. 2003.

- D. G. Stearns, P. Mirkarimi and E. Spiller. Localized defects in multilayer coatings. *Thin Solid Films*, 446(1):37–49, 2004.
- D. G. Stearns, M. B. Stearns, Y. Cheng, J. H. Stith and N. M. Ceglio. Thermally induced structural modification of Mo-Si multilayers. *Journal of Applied Physics*, 67(5):2415, 1990.
- G. Stengl, H. Löschner and P. Wolf. Ion projection lithography for sub-micron modification of materials. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 19-20:987–994, 1987.
- C. R. Stephens and H. Waelbroeck. Effective degrees of freedom in genetic algorithms. *Physical Review E*, 57(3):3251, 1998.
- C. R. Stephens, H. Waelbroeck and R. Aguirre. Schemata as Building Blocks: Does Size Matter? In *Foundations of Genetic Algorithms 5*, page 17. 2000.
- J. Stirniman and M. Rieger. Fast proximity correction with zone sampling. In *Proceedings of SPIE*, volume 2197, page 294. 1994.
- G. Stokes. On the composition and resolution of streams of polarized light from different sources. *Trans. Cambridge Philos. Soc.*, 9:399–416, 1852.
- R. Storn and K. Price. Differential evolution—a simple and efficient heuristic for global optimization over continuous spaces. *Journal of global optimization*, 11(4):341–359, 1997.
- F. Streichert, G. Stein, H. Ulmer and A. Zell. A clustering based niching ea for multimodal search spaces. In *Artificial Evolution*, pages 293–304. Springer, 2004.
- F. Su, F. Zhu, Z. Yin, H. Yao, Q. Wang and W. Dong. New Crossover Operator of Genetic Algorithms for the TSP. In *2009 International Joint Conference on Computational Sciences and Optimization*, pages 666–669. IEEE, 2009.
- M. B. Subrahmanyam. An extension of the simplex method to constrained nonlinear optimization. *Journal of Optimization Theory and Applications*, 62(2):311–319, 1989.
- M. Sugawara, A. Chiba, H. Yamanashi, H. Oizumi and I. Nishiyama. Attenuated phase-shift mask for line patterns in EUV lithography. *Microelectronic Engineering*, 67-68:10–16, 2003.
- J. Suh and D. Van Gucht. Incorporating heuristic information into genetic search. In *Proceedings of the Second International Conference on Genetic Algorithms on Genetic algorithms and their application*, pages 100–107. L. Erlbaum Associates Inc., 1987.
- A. Suzuki, K. Saitoh and M. Yoshii. Multilevel imaging system realizing $k_1=0.3$ lithography. In L. Van den Hove (editor), *Optical Microlithography XII*, volume 3679, pages 396–407. SPIE, Santa Clara, CA, USA, 1999.
- E. Suzuki and S. Kobayashi. Sample-based Crowding Method for Multimodal Optimization in Continuous Domain. In *Proceedings of the 2005 IEEE Congress on Evolutionary Computation*, volume 2, pages 1867–1874. IEEE, 2005.
- M. Switkes, M. Rothschild, T. A. Shedd, H. B. Burnett and M. S. Yeung. Bubbles in immersion lithography. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 23(6):2409–2412, 2005.
- G. Syswerda. Uniform crossover in genetic algorithms. In *Proceedings of the third international conference on Genetic algorithms*, pages 2–9. Morgan Kaufmann Publishers Inc., 1989.
- G. Syswerda. Schedule optimization using genetic algorithms. *Handbook of genetic algorithms*, pages 332–349, 1991.

- A. Taflové and M. Brodwin. Numerical Solution of Steady-State Electromagnetic Scattering Problems Using the Time-Dependent Maxwell's Equations. *IEEE Transactions on Microwave Theory and Techniques*, 23(8):623–630, 1975.
- E. Talbi, S. Mostaghim and T. Okabe. Parallel approaches for multiobjective optimization. In J. Branke, K. Deb, K. Miettinen and R. Słowiński (editors), *Multiobjective Optimization*, pages 349–372. Springer-Verlag Berlin, Heidelberg, 2008.
- S. Tanaka, H. Nakamura, K. Kawano and S. Inoue. Practical topography design for alternating phase-shifting mask. In *Proceedings of SPIE*, volume 2726, pages 473–484. SPIE, 1996.
- T. Tanaka. *Detection signal analysis of actinic inspection of EUV mask blanks using dark-field imaging*. SPIE, 2006.
- R. Tanese. *Distributed genetic algorithms for function optimization*. Ph.D. dissertation, University of Michiga, 1989.
- K. Tang, X. Li, P. N. Suganthan, Z. Yang and T. Weise. Benchmark Functions for the CEC ' 2010 Special Session and Competition on Large-Scale Global Optimization. *Optimization*, (1):1–23, 2010.
- R. Tapia. On secant updates for use in general constrained optimization. *Mathematics of Computation*, 51(183):181–202, 1988.
- S. Tarutani, H. Tsubaki and S. Kanna. Development of materials and processes for double patterning toward 32-nm node 193-nm immersion lithography process. In *Proceedings of SPIE*, volume 6923, page 69230F. 2008.
- D. M. Tate and A. E. Smith. Expected allele coverage and the role of mutation in genetic algorithms. In *Proceedings of the Fifth International Conference on Genetic Algorithms*, pages 31–37. 1993.
- J. C. Taylor, C. R. Chambers, R. Deschner, R. J. LeSuer, W. E. Conley, S. D. Burns and C. G. Willson. Implications of immersion lithography on 193-nm photoresists. In *Advances in Resist Technology and Processing XXI*, volume 5376, pages 34–43. SPIE, 2004.
- A. Teller. Evolving Programmers: The Co-evolution of Intelligent Recombination Operators. In P. J. Angeline and K. E. Kinneer JR (editors), *Advances in Genetic Programming 2*, volume 2, chapter 3, pages 45–68. MIT Press, 1996.
- L. Temby, P. Vamplew and A. Berry. Accelerating real-valued genetic algorithms using mutation-with-momentum. *AI 2005: Advances in Artificial Intelligence*, pages 1108–1111, 2005.
- Y. Teramoto, Z. Narihiro, D. Yamatani, T. Yokoyama, K. Bessho, Y. Joshima, T. Shirai, S. Mouri, T. Inoue, H. Mizokoshi and Others. Development of Sn-fueled high-power DPP EUV source for enabling HVM. In *Proceedings of SPIE*, volume 6517, pages 722–736. 2007.
- T. Terasawa, N. Hasegawa, H. Fukuda and S. Katagiri. Imaging Characteristics of Multi-Phase-Shifting and Halftone Phase-Shifting Masks. *Japanese Journal of Applied Physics*, 30(Part 1, No. 11B):2991–2997, 1991.
- T. Terasawa, N. Hasegawa, T. Kurosaki and T. Tanaka. 0.3-micron optical lithography using a phase-shifting mask. In *Proceedings of SPIE*, volume 1088, page 25. SPIE, 1989.
- Y. Tezuka, M. Ito, T. Terasawa and T. Tomie. Actinic detection of multilayer defects on EUV mask blanks using LPP light source and dark-field imaging. In *Proceedings of SPIE*, volume 5374, page 271. SPIE, 2004.
- M. Thiel, J. Fischer, G. von Freymann and M. Wegener. Direct laser writing of three-dimensional submicron structures using a continuous-wave laser at 532 nm. *Applied Physics Letters*, 97(22):221102, 2010.

- L. Thiele, K. Miettinen, P. J. Korhonen and J. Molina. A preference-based evolutionary algorithm for multi-objective optimization. *Evolutionary computation*, 17(3):411–36, 2009.
- D. Thierens. Dimensional Analysis of Allele-Wise Mixing Revisited. In *Proceedings of the 4th International Conference on Parallel Problem Solving from Nature*, PPSN IV, pages 255–265. Springer-Verlag, London, UK, 1996.
- D. Thierens and D. E. Goldberg. Convergence models of genetic algorithm selection schemes. In *Parallel Problem Solving from Nature III*, pages 119–129. Springer, 1994.
- G. Timp, R. Behringer, D. Tennant, J. Cunningham, M. Prentiss and K. Berggren. Using light as a lens for submicron, neutral-atom lithography. *Physical review letters*, 69(11):1636–1639, 1992.
- J. Tirapu-Azpiroz and P. Burchard. Boundary layer model to account for thick mask effects in photolithography. In *Proceedings of SPIE*, volume 5040, pages 1611–1619. SPIE, 2003.
- K. Toh, G. Dao, R. Singh and H. Gaw. Optical lithography with chromeless phase-shifted masks. In *Proceedings of SPIE*, volume 1463, page 74. SPIE, 1991a.
- K. K. Toh, G. T. Dao, R. R. Singh and H. T. Gaw. Chromeless phase-shifted masks: a new approach to phase-shifting masks. In *Proceedings of SPIE*, volume 1496, pages 27–53. SPIE, 1991b.
- B. Tollkühn, A. Erdmann, J. Lammers, C. Nölscher and A. Semmler. Do we need complex resist models for predictive simulation of lithographic process performance? In *Proceedings of SPIE*, volume 5376, pages 983–994. SPIE, 2004.
- B. Tollkühn, T. Fühner, D. Matiut, A. Erdmann, A. Semmler, B. Küchler and G. Kókai. Will Darwin’s law help us to improve our resist models? In *Proceedings of SPIE*, volume 5039, pages 291–302. SPIE, 2003.
- V. J. Torczon. *Multi-directional search: a direct search algorithm for parallel machines*. Ph.D. dissertation, Rice University Houston, TX, 1989.
- V. J. Torczon. On the Convergence of Pattern Search Algorithms. *SIAM Journal on Optimization*, 7(1):1–25, 1997.
- A. Törn and A. Žilinskas. *Global Optimization*, volume 350 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, Berlin, Heidelberg, 1989.
- Y. Trouiller, J. Serrand, C. Miramond, Y. Rody, S. Manakli and P. J. Goirand. ArF Imaging with Off-Axis Illumination and Subresolution Assist Bars, A Compromise between Mask Constraints and Lithographic Process Constraints. In *Proc. SPIE 4691*, page 1522. 2002.
- M.-F. Tsai, S.-J. Chang, C. C. P. Chen and L. S. Melvin III. Abbe-PCA (Abbe-Hopkins): microlithography aerial image analytical compact kernel generation based on principle component analysis. In *Proceedings of SPIE*, volume 7274, pages 72742D–72742D–9. SPIE, 2009.
- S. Tsutsui and D. E. Goldberg. Search space boundary extension method in real-coded genetic algorithms. *Information Sciences*, 133(3-4):229–247, 2001.
- P. Turney. Myths and Legends of the Baldwin Effect. *Proceedings of the ICML96 13th International Conference on Machine Learning*, pages 135–142, 1996.
- J. Turunen. Form-birefringence limits of Fourier-expansion methods in grating theory. *Journal of the Optical Society of America A*, 13(5):1013, 1996.
- W. Ulrich, H. J. Rostalski and R. Hudyma. Development of dioptric projection lenses for deep ultraviolet lithography at Carl Zeiss. *Journal of Microlithography, Microfabrication, and Microsystems*, 3:87, 2004.

- H. P. Urbach and D. A. Bernard. Modeling latent-image formation in photolithography, using the Helmholtz equation. *Journal of the Optical Society of America A*, 6(9):1343, 1989.
- R. Ursem. Multinational GAs: Multimodal optimization techniques in dynamic environments. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 19–26. 2000.
- J. Ushioda, Y. Seki, K. Maeda, T. Ohfuji and H. Tanabe. Chromium Fluoride Attenuated Phase-Shifting Mask for Argon Fluoride Excimer Laser Lithography. *Japanese Journal of Applied Physics*, 35(Part 1, No. 12B):6356–6359, 1996.
- Å. E. Uyar and A. E. Harmanci. Investigation of new operators for a diploid genetic algorithm. In *Proceedings of SPIE*, volume 3812, pages 32–43. SPIE, 1999.
- D. A. Van Steenwinckel and J. Lammers. Enhanced processing: sub-50 nm features with 0.8- μ m DOF using a binary reticle. In *Proceedings of SPIE*, pages 225–239. SPIE, 2003.
- D. A. Van Veldhuizen. *Multiobjective evolutionary algorithms: classifications, analyses, and new innovations*. Ph.D. dissertation, Air Force Institute of Technology, Wright-Patterson AFB, Ohio, 1999.
- D. A. Van Veldhuizen and G. Lamont. Multiobjective evolutionary algorithm research: A history and analysis. Technical report, Air Force Institute of Technology, Dayton, OH, 1998.
- D. A. Van Veldhuizen and G. B. Lamont. Multiobjective optimization with messy genetic algorithms. In *Proceedings of the 2000 ACM symposium on Applied computing - SAC '00*, pages 470–476. ACM Press, New York, New York, USA, 2000.
- A. Vardi. A trust region algorithm for equality constrained minimization: convergence properties and implementation. *SIAM Journal on Numerical Analysis*, 22(3):575–591, 1985.
- G. A. Vignaux and Z. Michalewicz. A genetic algorithm for the linear transportation problem. *IEEE Transactions on Systems, Man, and Cybernetics*, 21(2):445–452, 1991.
- G. M. Viswanathan, S. V. Buldyrev, S. Havlin, M. G. da Luz, E. P. Raposo and H. E. Stanley. Optimizing the success of random searches. *Nature*, 401(6756):911–4, 1999.
- H.-M. Voigt and T. Anheyer. Modal mutations in evolutionary algorithms. In *Evolutionary Computation, 1994. IEEE World Congress on Computational Intelligence., Proceedings of the First IEEE Conference on*, pages 88–92. IEEE, 1994.
- H.-M. Voigt, H. Mühlenbein and D. Cvetkovic. Fuzzy Recombination for the Breeder Genetic Algorithm. pages 104–113, 1995.
- R. M. von Büнау. Optimal coherent decompositions for radially symmetric optical systems. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 15(6):2412, 1997.
- M. D. Vose. Generalizing the notion of schema in genetic algorithms. *Artificial Intelligence*, 50(3):385–396, 1991.
- M. D. Vose. A Critical Examination Of The Schema Theorem. Technical report, University of Tennessee, 1993.
- M. D. Vose. *The Simple Genetic Algorithm: Foundations and Theory*. MIT Press, Cambridge, MA, USA, 1999.
- M. D. Vose and G. E. Liepins. Punctuated equilibria in genetic search. *Complex systems*, 5:31–44, 1991.
- M. D. Vose and A. H. Wright. Form invariance and implicit parallelism. *Evolutionary computation*, 9(3):355–70, 2001.

- C. Wagner and N. Harned. EUV lithography: Lithography gets extreme. *Nature Photonics*, 4(1):24–26, 2010.
- H. Watanabe, Y. Todokoro and M. Inoue. Transparent phase shifting mask. In *Electron Devices Meeting, 1990. IEDM'90. Technical Digest., International*, pages 821–824. IEEE, 1990.
- X. Wei, A. J. Wachters and H. P. Urbach. Finite-element model for three-dimensional optical scattering problems. *Journal of the Optical Society of America A*, 24(3):866, 2007.
- T. Weiland. A discretization model for the solution of Maxwell's equations for six-component fields. *Archiv Elektronik und Übertragungstechnik*, 31(3):116–120, 1977.
- R. Weinberg and M. Berkus. Computer simulation of a living cell: Part I. *International Journal of Bio-Medical Computing*, 2(2):95–120, 1971.
- M. Weiß, H. Binder and R. Schwalm. Modeling and simulation of a chemically amplified DUV resist using the effective acid concept. *Microelectronic Engineering*, 27(1-4):405–408, 1995.
- P. Werbos. Backpropagation through time: what it does and how to do it. *Proceedings of the IEEE*, 78(10):1550–1560, 1990.
- T. White and F. Oppacher. Adaptive crossover using automata. pages 229–238. Springer, 1994.
- L. D. Whitley, S. Rana, J. Dzubera and K. Mathias. Evaluating Evolutionary Algorithms. *Artificial Intelligence*, 85(1-2):245–276, 1996.
- L. D. Whitley. An Executable Model of a Simple Genetic Algorithm. In D. Whitley (editor), *Foundations of Genetic Algorithms Workshop FOGA92*, pages 45–62. Morgan Kaufmann, 1992.
- L. D. Whitley. A Free Lunch Proof for Gray versus Binary Encodings. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 726–733. Morgan Kaufmann, San Francisco, 1999.
- L. D. Whitley and K. Kauth. GENITOR: A different genetic algorithm. In *Proceedings of the 1988 Rocky Mountain Conference on Artificial Intelligence*, pages 118–130. 1988.
- L. D. Whitley, K. Mathias and P. Fitzhorn. Delta-Coding: An Iterative Search Strategy for Genetic Algorithms. Morgan Kaufmann, 1991.
- L. D. Whitley, T. Starkweather and D. Fuquay. Scheduling Problems and Traveling Salesmen: The Genetic Edge Recombination Operator. In *Proceedings of the 3rd International Conference on Genetic Algorithms*, pages 133–140. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1989.
- M. J. Wieland, G. de Boer, G. F. ten Berge, M. van Kervinck, R. Jager, J. J. M. Peijster, E. Slot, S. W. H. K. Steenbrink, T. F. Teepen and B. J. Kampherbeek. MAPPER: high-throughput maskless lithography. In *Proceedings of SPIE*, pages 76370F–76370F–10. 2010.
- O. Wiener. Stehende Lichtwellen und die Schwingungsrichtung polarisierten Lichtes. *Annalen der Physik*, 276(6):203–243, 1890.
- A. P. Wierzbicki. *The use of reference objectives in multiobjective optimization*. Lecture notes in economics and mathematical systems. Springer-Verlag, 1980.
- R. B. Wilson. *A simplicial algorithm for concave programming*. Ph.D. dissertation, Harvard University, 1965.
- S. W. Wilson. *Classifier system learning of a boolean function*. Technical report, Rowland Institute For Science, 1986.
- G. Wojcik, D. Vaughan and L. Galbraith. Calculation of light scatter from structures on silicon surfaces. In *Proceedings of SPIE*, volume 774. 1987.

- D. Wolpert and W. Macready. No free lunch theorems for optimization. *IEEE Transactions on Evolutionary Computation*, 1(1):67–82, 1997.
- D. Wolpert and W. Macready. Coevolutionary Free Lunches. *IEEE Transactions on Evolutionary Computation*, 9(6):721–735, 2005.
- A. K.-K. Wong. *Resolution enhancement techniques in optical lithography*. SPIE Press, Bellingham, WA, 2001.
- A. K.-K. Wong. *Optical imaging in projection microlithography*. SPIE Press, 2005.
- A. K.-K. Wong, R. A. Ferguson and S. M. Mansfield. The mask error factor in optical lithography. *IEEE Transactions on Semiconductor Manufacturing*, 13(2):235–242, 2000.
- A. K.-K. Wong, L. W. Liebmann and A. Molless. Alternating phase-shifting mask with reduced aberration sensitivity: lithography considerations. In *Proceedings of SPIE*, volume 4346, pages 420–428. SPIE, 2001.
- A. K.-K. Wong and A. Neureuther. Mask topography effects in projection printing of phase-shifting masks. In *IEEE Transactions on Electron Devices*, volume 41, pages 895–902. 1994.
- A. K.-K. Wong and A. R. Neureuther. Edge effects in phase-shifting masks for 0.25- μm lithography. In *Proceedings of SPIE*, volume 222, pages 222–228. SPIE, 1993.
- A. K.-K. Wong and A. R. Neureuther. Rigorous three-dimensional time-domain finite-difference electromagnetic simulation for photolithographic applications. *IEEE Transactions on Semiconductor Manufacturing*, 8(4):419–431, 1995.
- D. J. Woods. *An interactive approach for solving multi-objective optimization problems*. Ph.D. dissertation, Rice University Houston, TX, 1985.
- A. H. Wright. Genetic algorithms for real parameter optimization. *Foundations of genetic algorithms*, 1:205–218, 1991.
- A. H. Wright. The Exact Schema Theorem, 1999.
- A. H. Wright, M. D. Vose and J. E. Rowe. Implicit Parallelism. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 1505–1517. Springer, 2003.
- E. Wu, J. Strickler, W. Harrell and W. Webb. Two-photon lithography for microelectronic application. In *Proceedings of SPIE*, volume 1674, pages 776–782. SPIE, 1992.
- S. Wurm, H. Han, P. Kearney, W. Cho, C. Jeon and E. Gullikson. EUV mask blank defect inspection strategies for 32-nm half-pitch and beyond. In *Proceedings of SPIE*, volume 6607, page 66073A. 2007.
- X. S. Xie, M. Li, J. Guo, B. Liang, Z. X. Wang, A. Sinitskii, Y. Xiang and J. Y. Zhou. Phase manipulated multi-beam holographic lithography for tunable optical lattices. *Optics Express*, 15(11):7032, 2007.
- E. Yablonovitch and R. B. Vrijen. Optical projection lithography at half the Rayleigh resolution limit by two-photon exposure. *Optical Engineering*, 38(2):334, 1999.
- A. Yamada, Y. Oae, T. Okawa, M. Takizawa and M. Yamabe. Evaluation of throughput improvement by MCC and CP in multicolumn e-beam exposure system. In *Proceedings of SPIE*, volume 7637, page 76370C. 2010.
- N. Yamamoto, J. Kye and H. J. Levinson. Mask topography effect with polarization at hyper NA. In *Proceedings of SPIE*, volume 6154, page 61544F. 2006.
- K. Yamazoe. Computation theory of partially coherent imaging by stacked pupil shift matrix. *Journal of the Optical Society of America A*, 25(12):3111–3119, 2008.

- P. Y. Yan. Understanding Bossung curve asymmetry and focus shift effect in EUV lithography. In *Proceedings of SPIE*, volume 4562, page 279. 2002.
- P. Y. Yan, G. Zhang, P. Kofron, J. Chow, A. Stivers, E. Tejnli, G. Cardinale and P. Kearney. EUV mask patterning approaches. In *Proceedings of SPIE*, volume 3676, page 309. 1999.
- P. Y. Yan, S. Yan, G. Zhang, P. Kearney, J. Richards, P. Kofron and J. Chow. EUV mask absorber defect repair with focused ion beam. In *Proceedings of SPIE*, volume 3546, page 206. SPIE, 1998.
- Y. Yanagishita, N. Ishiwata, Y. Tabata, K. H. Nakagawa and K. Shigematsu. Phase-shifting photolithography applicable to real IC patterns. In *Proceedings of SPIE*, volume 1463, page 207. SPIE, 1991.
- M. Yanagiya. A simple mutation-dependent genetic algorithm. In *Proceedings of the 5th International Conference on Genetic Algorithms*, page 659. Morgan Kaufmann Publishers Inc., 1993.
- X.-S. Yang. Firefly algorithm, stochastic test functions and design optimisation. *International Journal of Bio-Inspired Computation*, 2(2):78–84, 2010.
- X.-S. Yang. Firefly Algorithms for Multimodal Optimization. In *Stochastic Algorithms: Foundations and Applications: 5th International Symposium*, volume 5792, pages 169–178. Springer-Verlag New York, 2009.
- X.-S. Yang and S. Deb. Cuckoo Search via Lévy flights. In *World Congress on Nature & Biologically Inspired Computing, 2009. NaBIC 2009.*, pages 210–214. IEEE, 2009.
- X. Yao. A review of evolutionary artificial neural networks. *International Journal of Intelligent Systems*, 3(1):539–567, 1993.
- K. S. Yee. Numerical solution of initial boundary value problems involving Maxwell’s equations in isotropic media. *IEEE Transactions on Antennas and Propagation*, 14(3):302–307, 1966.
- M. S. Yeung. Modeling aerial images in two and three dimensions. *Proc. Kodak Microelectronics Seminar*, 1985.
- M. S. Yeung. Modeling high numerical aperture optical lithography. In *Proceedings of SPIE*, volume 922, pages 149–167. 1988.
- M. S. Yeung and E. Barouch. *Three-dimensional nonplanar lithography simulation using a periodic fast multipole method*. SPIE, 1997.
- M. S. Yeung, D. Lee, R. S. Lee and A. R. Neureuther. Extension of the Hopkins theory of partially coherent imaging to include thin-film interference effects. In *Proceedings of SPIE*, volume 1927, pages 452–463. 1993.
- X. Yin and N. Gernay. A fast genetic algorithm with sharing scheme using cluster analysis methods in multimodal function optimization. In *Artificial Neural Nets and Genetic Algorithms*, pages 450–457. Springer, 1993.
- Y. Yoshida and N. Adachi. A Diploid Genetic Algorithm for Preserving Population Diversity - pseudo-Meiosis GA. In *The Third Conference on Parallel Problem Solving from Nature*, pages 36–45. Springer, 1994.
- M. Yoshizawa, V. Philipsen, L. Leunissen, E. Hendrickx, R. Jonckheere, G. Vandenberghe, U. Buttgerit, H. Becker, C. Koepf and M. Irmscher. Comparative study of bi-layer attenuating phase-shifting masks for hyper-NA lithography. In *Proceedings of SPIE*, volume 6283, page 62831G. 2006.
- P. L. Yu. A Class of Solutions for Group Decision Problems. *Management Science*, 19(8):936–946, 1973.

- P. L. Yu. Cone convexity, cone extreme points, and nondominated solutions in decision problems with multiobjectives. *Journal of Optimization Theory and Applications*, 14(3):319–377, 1974.
- C.-M. Yuana and V. Polb. Comparisons of attenuated PSM technologies using etched quartz and embedded materials. *Microelectronic Engineering*, 27(1-4):251–253, 1995.
- S. Yulin, N. Benoit, T. Feigl and N. Kaiser. Interface-engineered EUV multilayer mirrors. *Microelectronic Engineering*, 83(4-9):692–694, 2006.
- L. Zadeh. Optimality and non-scalar-valued performance criteria. *IEEE Transactions on Automatic Control*, 8(1):59–60, 1963.
- D. Zaharie. A comparative analysis of crossover variants in differential evolution. *Proceedings of the IMCSIT*, pages 183–192, 2007.
- S. Zegordi, I. K. Abadi and M. B. Nia. A novel genetic algorithm for solving production and transportation scheduling in a two-stage supply chain. *Computers & Industrial Engineering*, 58(3):373–381, 2010.
- F. Zernike. Beugungstheorie des Schneidenverfahrens und seiner verbesserten Form, der Phasenkontrastmethode. *Physica*, 1(7-12):689–704, 1934.
- F. Zernike. The concept of degree of coherence and its application to optical problems. *Physica*, 5(8):785–795, 1938.
- J. Zhang, W. Xiong, Y. Wang, Z. Yu and M.-C. Tsai. A highly efficient optimization algorithm for pixel manipulation in inverse lithography technique. In *Proceedings of the 2008 IEEE/ACM International Conference on Computer-Aided Design, ICCAD '08*, pages 480–487. IEEE Press, Piscataway, NJ, USA, 2008.
- M. Zhang, S. Zhao and X. Wang. A hybrid self-adaptive genetic algorithm based on sexual reproduction and Baldwin effect for global optimization. In *Proceedings of the Eleventh Congress on Evolutionary Computation, CEC'09*, pages 3087–3094. IEEE Press, Piscataway, NJ, USA, 2009.
- J. Zhou, N. V. Lafferty, B. W. Smith and J. H. Burnett. Immersion lithography with numerical apertures above 2.0 using high index optical materials. In *Optical Microlithography XX*, volume 6520, page 65204T. SPIE, 2007.
- P. A. Zimmerman, C. V. Peski, B. Rice, J. D. Byers, N. J. Turro, X. Lei, J. L. Gejo, V. Liberman, S. Palmacci, M. Rothchild, A. Whitker, I. Blakey, L. Chen, B. Dargaville and H. Liu. Status of High-Index Materials for Generation-Three 193nm Immersion Lithography. *Journal of Photopolymer Science and Technology*, 20(5):643–650, 2007.
- E. Zitzler. *Evolutionary algorithms for multiobjective optimization: Methods and applications*. Ph.D. dissertation, Swiss Federal Institute of Technology (ETH) Zurich, Zurich, Switzerland, 1999.
- E. Zitzler, K. Deb and L. Thiele. Comparison of multiobjective evolutionary algorithms: empirical results. *Evolutionary computation*, 8(2):173–95, 2000.
- E. Zitzler and S. Künzli. Indicator-based selection in multiobjective search. In *Parallel Problem Solving from Nature (PPSN VIII)*, Lecture Notes in Computer Science, pages 832–842. Springer, Heidelberg, Germany, 2004.
- E. Zitzler, M. Laumanns and S. Bleuler. A tutorial on evolutionary multiobjective optimization. In *Metaheuristics for Multiobjective Optimisation*, pages 3–37. Springer, 2004.
- E. Zitzler, M. Laumanns and L. Thiele. SPEA2: Improving the strength Pareto evolutionary algorithm. Technical Report 103, Swiss Federal Institute of Technology (ETH), Zurich, Switzerland, 2001.

- E. Zitzler and L. Thiele. An evolutionary algorithm for multiobjective optimization: The strength Pareto approach. Technical Report 43, Swiss Federal Institute of Technology (ETH), Zurich, Switzerland, 1998a.
- E. Zitzler and L. Thiele. Multiobjective optimization using evolutionary algorithms—a comparative case study. In *Parallel Problem Solving from Nature—PPSN V*, pages 292–301. Springer, 1998b.
- E. Zitzler, L. Thiele, M. Laumanns, C. Fonseca and V. Da Fonseca. Performance assessment of multiobjective optimizers: An analysis and review. *IEEE Transactions on Evolutionary Computation*, 7(2):117–132, 2003.
- J. Zydallis. *Explicit building-block multiobjective genetic algorithms: Theory, analysis, and development*. Ph.D. dissertation, Air University, Wright-Patterson, Air force Base, Ohio, U.S.A., 2003.

List of Own Publications

- T. Fühner and C. Jacob. EvolVision - an Evolvica visualization tool. In L. Spector, E. D. Goodman, A. Wu, W. B. Langdon, H.-M. Voigt, M. Gen, S. Sen, M. Dorigo, S. Pezeshk, M. H. Garzon and E. Burke (editors), *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2001)*, page 176. Morgan Kaufmann, San Francisco, California, USA, 2001.
- A. Erdmann, R. Farkas, T. Fühner, B. Tollkühn and G. Kókai. Mask and source optimization for lithographic imaging systems. In F. Wyrowski (editor), *Proceedings of SPIE*, volume 5182, pages 88–102. SPIE, 2003.
- B. Tollkühn, T. Fühner, D. Matiut, A. Erdmann, A. Semmler, B. Kuchler and G. Kókai. Will Darwin’s law help us to improve our resist models? In *Proceedings of SPIE*, volume 5039, pages 291–302. SPIE, 2003.
- A. Erdmann, T. Fühner, T. Schnattinger and B. Tollkühn. Towards Automatic Mask and Source Optimization for Optical Lithography. In *Proceedings of SPIE*, volume 5377, pages 646–657. SPIE, 2004.
- T. Fühner, A. Erdmann, R. Farkas, B. Tollkühn and G. Kókai. Genetic algorithms to improve mask and illumination geometries in lithographic imaging systems. In R. et Al. (editor), *Applications of Evolutionary Computing*, volume 3005, pages 208–218. Springer, 2004a.
- T. Fühner, A. Erdmann, C. J. Ortiz and J. Lorenz. Genetic algorithm for optimization and calibration in process simulation. *Simulation of Semiconductor ...*, page 347, 2004b.
- T. Fühner, A. Erdmann and T. Schnattinger. Genetic algorithms for geometry optimization in lithographic imaging systems. In A. G. Tescher (editor), *Proceedings of SPIE*, volume 5558, pages 29–40. SPIE, 2004c.
- T. Fühner and T. Jung. Use of genetic algorithms for the development and optimization of crystal growth processes. *Journal of Crystal Growth*, 266(1-3):229–238, 2004.
- C. J. Ortiz, P. Pichler, T. Fühner, F. Cristiano, B. Colombeau, N. E. B. Cown and A. Claverie. A physically based model for the spatial and temporal evolution of self-interstitial agglomerates in ion-implanted silicon. *Journal of Applied Physics*, 96(9):4866–4877, 2004.
- B. Fischer, J. Friedrich, T. Jung, M. Hainke, J. Dagner, T. Fühner and P. Schwesig. Modeling of industrial bulk crystal growth—state of the art and challenges. *Journal of Crystal Growth*, 275(1-2):240–250, 2005.
- T. Fühner and A. Erdmann. Improved mask and source representations for automatic optimization of lithographic process conditions using a genetic algorithm. In B. W. Smith (editor), *Proceedings of SPIE*, volume 5754, pages 415–426. SPIE, 2005.
- A. Erdmann, P. Evanschitzky, G. Citarella, T. Fühner and P. De Bisschop. Rigorous mask modeling using waveguide and FDTD methods: an assessment for typical hyper-NA imaging problems. In *Proceedings of SPIE*, volume 6283, page 628319. 2006.
- T. Fühner, S. Popp, C. Dürr and A. Erdmann. Efficient optimization of lithographic process conditions using a distributed combined global/local search approach. In D. G. Flagello (editor), *Proceedings of SPIE*, volume 61543S, pages 61543S–61543S–12. SPIE, 2006a.

- T. Fühner, S. Popp and T. Jung. A novel framework for distributing computations DisPyTE – distributing Python tasks environment. *Journal of Computational Electronics*, 5(4):349–352, 2006b.
- C. Dürr, T. Fühner and P. Suganthan. LisBON: A framework for parallelisation and hybridisation of optimisation algorithms. In *2007 IEEE Congress on Evolutionary Computation*, pages 1717–1724. IEEE, 2007a.
- C. Dürr, T. Fühner, B. Tollkühn, A. Erdmann and G. Kókai. Memetic Algorithms: Parametric Optimization for Microlithography. In *Hybrid Evolutionary Algorithms*, pages 201–239. Springer, 2007b.
- A. Erdmann, T. Fühner, S. Seifert, S. Popp and P. Evanschitzky. The impact of the mask stack and its optical parameters on the imaging performance. In D. G. Flagello (editor), *Proceedings of SPIE*, volume 6520, pages 65201I–65201I–12. SPIE, 2007.
- T. Fühner, A. Erdmann and S. Seifert. Direct optimization approach for lithographic process conditions. *Journal of Micro/Nanolithography, MEMS and MOEMS*, 6(3):031006, 2007a.
- T. Fühner, T. Schnattinger, G. Ardelean and A. Erdmann. Dr.LiTHO: a development and research lithography simulator. In *Proceedings of SPIE*, page 65203F. SPIE, 2007b.
- A. Erdmann, P. Evanschitzky, T. Fühner, T. Schnattinger, C. Xu and C. Szmanda. Rigorous electromagnetic field simulation of two-beam interference exposures for the exploration of double patterning and double exposure scenarios. In *Proceedings of SPIE*, volume 6924, page 692452. 2008a.
- A. Erdmann, T. Fühner and P. Evanschitzky. Optimization of mask absorber stacks and illumination settings for contact hole imaging. In *Proceedings of SPIE*, volume 7028, pages 70283L–70283L–12. SPIE, 2008b.
- A. Erdmann, T. Fühner and P. Evanschitzky. Optimization of mask absorber stacks and illumination settings for contact hole imaging. In T. Horiuchi (editor), *Proceedings of SPIE*, volume 7028, pages 70283L–70283L–12. 2008c.
- A. Erdmann, D. Reibold, T. Fühner and P. Evanschitzky. Photomasks for Semiconductor Lithography: From Simple Shadow Casters to Complex 3D Scattering Objects. *Advances in Science and Technology*, 55:173–180, 2008d.
- T. Fühner, A. Erdmann and P. Evanschitzky. Simulation-based EUV source and mask optimization. In *Proceedings of SPIE*, volume 7122, pages 71221Y–71221Y–14. SPIE, 2008a.
- T. Fühner, C. Kampen, I. Kodrasi, A. Burenkov and A. Erdmann. A simulation study on the impact of lithographic process variations on CMOS device performance. In H. J. Levinson and M. V. Dusa (editors), *Proceedings of SPIE*, volume 6924 of *Presented at the Society of Photo-Optical Instrumentation Engineers (SPIE) Conference*, pages 692453–692453–12. SPIE, 2008b.
- C. Kampen, T. Fühner, A. Burenkov, A. Erdmann, J. Lorenz and H. Ryssel. On the stability of fully depleted SOI MOSFETs under lithography process variations. In *ESSDERC 2008 - 38th European Solid-State Device Research Conference*, pages 194–197. IEEE, 2008.
- B. Meliorisz, S. Partel, T. Schnattinger, T. Fühner, A. Erdmann and P. Hudek. Investigation of high-resolution contact printing. *Microelectronic Engineering*, 85(5-6):744–748, 2008.
- A. Erdmann, P. Evanschitzky and T. Fühner. Mask diffraction analysis and optimization for EUV masks. In *Proceedings of SPIE*, volume 7271, page 72711E. 2009a.
- A. Erdmann, T. Fühner, F. Shao and P. Evanschitzky. Lithography simulation: modeling techniques and selected applications. In H. Bosse, B. Bodermann and R. M. Silver (editors), *Proceedings of SPIE*, pages 739002–739002–17. SPIE, 2009b.

- P. Evanschitzky, A. Erdmann and T. Fühner. Extended Abbe approach for fast and accurate lithography imaging simulations. In U. F. W. Behringer (editor), *Proceedings of SPIE*, volume 7470, pages 747007–747007–11. SPIE, 2009.
- J. Lorenz, C. Kampen, A. Burenkov and T. Fühner. Impact of lithography variations on advanced CMOS devices. In *2009 International Symposium on VLSI Technology, Systems, and Applications*, pages 17–18. IEEE, 2009.
- J. S. Petersen, R. T. Greenway, T. Fühner, P. Evanschitzky, F. Shao and A. Erdmann. Exploration of linear and non-linear double exposure techniques by simulation. In *Proceedings of SPIE*, volume 7274, page 72741V. SPIE, 2009.
- F. Shao, P. Evanschitzky, T. Fühner and A. Erdmann. Rigorous diffraction simulations of topographic wafer stacks in double patterning. *Microelectronic Engineering*, 86(4-6):489–491, 2009a.
- F. Shao, P. Evanschitzky, T. Fühner and A. Erdmann. Efficient analysis of three dimensional EUV mask induced imaging artifacts using the waveguide decomposition method. In L. S. Zurbrick and M. W. Montgomery (editors), *Proceedings of SPIE*, volume 7488, pages 74882C–74882C–9. SPIE, 2009b.
- F. Shao, P. Evanschitzky, T. Fühner and A. Erdmann. Efficient simulation and optimization of wafer topographies in double patterning. *Journal of Micro/Nanolithography, MEMS and MOEMS*, 8(4):043070, 2009c.
- A. Erdmann, P. Evanschitzky and T. Fühner. Mask diffraction analysis and optimization for extreme ultraviolet masks. *Journal of Micro/Nanolithography, MEMS and MOEMS*, 9(1):13005, 2010a.
- A. Erdmann, F. Shao, P. Evanschitzky and T. Fühner. Mask-topography-induced phase effects and wave aberrations in optical and extreme ultraviolet lithography. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 28(6):C6J1, 2010b.
- A. Erdmann, F. Shao, P. Evanschitzky and T. Fühner. Mask and Wafer Topography Effects in Optical and EUV-Lithography. In *ECS Transactions*, volume 27, pages 415–420. 2010c.
- P. Evanschitzky, T. Fühner, F. Shao and A. Erdmann. Efficient simulation of three-dimensional EUV masks for rigorous source mask optimization and mask induced imaging artifact analysis. In U. F. Behringer and W. Maurer (editors), *26th European Mask and Lithography Conference*, volume 7545, pages 75450D–10. SPIE, Grenoble, France, 2010.
- S. Liu, T. Fühner, F. Shao, A. Barenbaum, J. Jahn and A. Erdmann. Topography-aware BARC optimization for double patterning. In M. V. Dusa and W. Conley (editors), *Proceedings of SPIE*, volume 7640, pages 76403C–76403C–12. 2010.
- V. Agudelo, P. Evanschitzky, A. Erdmann, T. Fühner, F. Shao, S. Limmer and D. Fey. Accuracy and performance of 3D mask models in optical projection lithography. In M. V. Dusa (editor), *Proceedings of SPIE*, volume 7973, pages 79730O–79730O–11. 2011.
- A. Erdmann, P. Evanschitzky, F. Shao, T. Fühner, G. F. Lorusso, E. Hendrickx, M. Goethals, R. Jonckheere, T. Bret and T. Hofmann. Predictive modeling of EUV-Lithography: the role of mask, optics, and photoresist effects. In D. G. Smith, F. Wyrowski and A. Erdmann (editors), *Proceedings of SPIE*, volume 8171, pages 81710M–81710M–16. 2011a.
- A. Erdmann, F. Shao, V. Agudelo, T. Fühner and P. Evanschitzky. Modeling of mask diffraction and projection imaging for advanced optical and EUV lithography. *Journal of Modern Optics*, 58(5-6):480–495, 2011b.

- P. Evanschitzky, T. Fühner and A. Erdmann. Image simulation of projection systems in photolithography. In B. Bodermann (editor), *Proceedings of SPIE*, volume 8083, pages 80830E–80830E–10. 2011a.
- P. Evanschitzky, F. Shao, T. Fühner and A. Erdmann. Compensation of mask induced aberrations by projector wavefront control. In M. V. Dusa (editor), *Proceedings of SPIE*, volume 7973, pages 797329–797329–8. 2011b.
- V. Agudelo, P. Evanschitzky, A. Erdmann and T. Fühner. Evaluation of various compact mask and imaging models for the efficient simulation of mask topography effects in immersion lithography. In W. Conley (editor), *Proceedings of SPIE*, volume 8326, pages 832609–832609–8. 2012.
- T. Fühner, P. Evanschitzky and A. Erdmann. Mutual source, mask and projector pupil optimization. In W. Conley (editor), *Proceedings of SPIE*, volume 8326, pages 83260I–83260I–12. 2012.
- V. Agudelo, T. Fühner, A. Erdmann and P. Evanschitzky. Application of artificial neural networks to compact mask models in optical lithography simulation. In W. Conley (editor), *Proceedings of SPIE*, volume 8683, pages 868317–868317–15. 2013.
- A. Erdmann, T. Fühner, P. Evanschitzky, J. T. Neumann, J. Ruoff and P. Gräupner. Modeling studies on alternative EUV mask concepts for higher NA. In P. P. Naulleau (editor), *Proceedings of SPIE*, volume 8679, pages 86791Q–86791Q–12. 2013.

Index

- 1/5th success rule, 81
- a posterior methods, 129, 131, 132
- a priori methods, 129, 131, 132
- aberration, *see* wavefront aberration
- Ackley's function, 175, 329
- activation function, 181
- active set, 162
- adaptive memetic algorithm, 150
- aerial image, 15, 16, 65
- Airy disk, 12
- Airy pattern, *see* Airy disk
- allele, 90, 122
- alphabet, 91
- alternating phase-shifting mask, *see*
 - photomask, alternating phase-shifting mask
- Ampere's law, 307
- ANN, *see* artificial neural network
- anti-reflective coating, 288
 - bottom anti-reflective coating, 71
 - top anti-reflective coating, 71
- anti-serif, 29
- archels, 207
- archive, 136, 139, 144
 - reduction, 144
- Argon fluoride laser, 2
- artificial immune system, 88, 99
- artificial neural network, 178, 179, 266
 - rescaling, 187
- attainment indicator, 147
- attainment method, 130
- attenuated phase shifting mask, 269
- attenuated phase-shifting mask, *see*
 - photomask, attenuated phase-shifting mask
- backpropagation learning, 184
- Baldwin effect, 97, 149
- Baldwinian evolution, 149
- Baldwinian learning, 149
- Band criterion, 237
- basins of attraction, 113
- batch learning, 185
- Beer's law, 74
- behavioral memory, 99
- Bessel function, 12
- BFGS approach, *see*
 - Broyden-Fletcher-Goldfarb-Shanno approach
- bi-objective optimization, 296
- binary tree, 197
- binary-coded genetic algorithm, 195, 229
- bit-string, 91
- Bossung curves, 16
- bottom anti-reflective coating, 288
 - bi-layer, 291
- boundary element method, 55
- boundary-layer method, 227
- Broyden-Fletcher-Goldfarb-Shanno approach,
 - 165, 186
- building block hypothesis, 120
- building blocks, 110, 120
- bulk image, 252
- catoptrics, 8
- CD criterion, 236
- cell removal method, 77
- chrome-on-glass mask, 206
- chromeless phase lithography, 225
- chromosome, 85, 90, 197
 - graph, 85
 - tree, 85
- Cittert-Zernike theorem, 315
- closing operator, 235
- coevolving memetic algorithm, 151
- coherence
 - partial, 17, 32, 60
 - spatial, 16
- coherent imaging, 313
- compromise programming, 131
- computational lithography, 205
- condenser, 8
- conflicting objectives, 125
- constitutive relations, 307
- constraints, 97, 232
- contact holes, 241, 269, 285
- contact lithography, 8
- continuous mask representation, 230
- contrast, 20

- contrast enhancement layer, 51
 - reversible, 51
- control parameters, 110
 - adaptation, 111
- convergence, 124, 136
- convergence rate, 162, 167
- corner rounding, 29
- CQuad illumination, 266, 272
- critical dimension, 10, 15
 - non-linearity, 28
- crossover, 83, 107
 - point, 85
 - two-point, 85
 - uniform, 85
- crossover rate, 110, 239
- crowding comparison operator, 139
- DAG, *see* directed acyclic graph
- deception, 91, 120
- decision maker, 129–131, 147
- decision space, 128
- decision space representation, 90
 - binary-coded genetic algorithms, 91
 - non-binary discrete alphabets, 92
 - real-coded genetic algorithms, 91
- decision variables, 128
- decoders, 99, 232
- deep-ultraviolet lithography, 7
- delay, 75
- density estimator, 139, 142
- depth of focus, 14, 15, 259
- develop, 7, 76
- Development and Research Lithography
 - simulator, 195
- development rate, 76
- di-allelic, 96
- differential evolution, 83
- diffraction, 8, 10, 12–14, 16, 309
 - efficiency, 43
 - Fraunhofer, 12, 17
 - Fraunhofer approximation, 312
 - Fresnel, 12
 - Fresnel approximation, 312
 - Fresnel-Kirchhoff, 309
 - grating, 16
 - Huygens-Fresnel principle, 310
 - Kirchhoff-Fresnel, 55
 - orders, 16, 32
 - Rayleigh-Sommerfeld, 55, 309
 - spectrum, 55, 62
 - incidence-angle-specific, 66
- diffraction order analysis, 207
- diffraction-limited, 10
- diffractive optical element, 36, 210
- diffusion, 74, 75
- dilation, 235
- Dill's model, 74
- Dill's parameters, 74
- diploidy, 96
- dipole illumination, 266
- direct laser writing, 26
- direct local variation method, 216
- direct source/mask optimization, 228
- direct write, 26
- directed acyclic graph, 138, 179
- discrete recombination, 82
- discretization, 156
- distributed Python tasks environment, 200
- diversification, 80
- diversity, 112
- diversity preservation, 112
- domain-decomposition techniques, 56, 57, 227
- dominance, 96
- double exposure, 41, 50
- double patterning, 51, 290
 - litho-cure-litho-etch, *see* double patterning, litho-freeze-litho-etch
 - litho-etch-litho-etch, 51
 - litho-freeze-litho-etch, 51
 - litho-litho-etch, 51
 - self-aligned, 52
 - space-defined, *see* double patterning, self-aligned
- down-hill simplex method, *see* Nelder-Mead method, *see* Nelder-Mead method
- Dr.LiTHO, *see* Development and Research Lithography simulator
- dynamic models, 123
- dynamic systems models, 124
- e-beam lithography, *see* electron-beam lithography
- effective acid concentration, 76
- electro-magnetic field effects, 227
- electro-magnetic field simulation, 252, 272, 291
- electro-magnetic field solver, 55
- electron-beam lithography, 26
- electronic design automation, 3
- elitism, 124, 136, 139
- elitist selection, 101, 197
- epistasis, 122, 230
- epsilon operator, 145
 - additive, 145
- epsilon-constraints genetic algorithm, 134
- epsilon-constraints method, 132
- equivalence class, 91

- erosion, 235
- event-driven network communication, 201
- evolution strategy, 80
- evolutionary algorithm, 79
- evolutionary algorithm control parameters
 - adaptation, 323
 - adaptive, 323
 - deterministic, 323
 - dynamic, 323
 - self-adaptive, 324
- evolutionary algorithms, 3
- evolutionary programming, 82
- exclusive-or gate, 180
- exploitation, 112, 134
- exploration, 112, 134
- exposure, 5, 74
- exposure latitude, 15, 259, 265
- extended source, 314
- eXtensible Mark-up Language remote
 - procedure call, 202
- extreme-ultraviolet lithography, 2, 8, 22, 278

- Faraday's law, 307
- fast marching method, 77
- feasibility-preserving methods, 99
- feasible objective space, 128
- feed-forward neural network, 179
- Fienup's algorithm, 214
- figures of merit, 232
- filling factor, *see* coherence, partial
- finite difference method, 75
- finite element method, 55
- finite integration technique, 55
- finite-difference time-domain, 55, 56
 - perfectly matched layers, 57
- first-order necessary condition, *see*
 - Karush-Kuhn-Tucker necessary condition
- fitness, 89
- fitness function, 152, 237
- fitness scaling, 90, 102
- focus shift, 278
- focused ion beam, 26
- forbidden pitch phenomenon, 241
- Fourier transform, 55
 - fast Fourier transform, 55
- frequency doubling, 39
- function approximation, 178

- g-line, 8
- GA-easy, 122
- GA-hard, 122
- Gauss's law, 307

- genders, 97
- gene, 90, 122
- gene expression genetic algorithm, 122
- genetic algorithm, 79, 228
- genetic algorithms theory, 118
- genetic drift, 109, 112
- genetic hitchhiking, 92
- genetic local search, 149
- genetic programming, 85
- genotype, 90, 149, 228
- Gerchberg-Saxton algorithm, 214
- GiNGA is a Niche Genetic Algorithm, 195
- global contrast, 265
- global convergence, 165
- global optimization, 227
- goal programming, 129, 131
- gradient descent method, 216
- gradient descent with momentum, 185
- gradient-based search approach, 218, 223
- granularity, 156
- Gray code, 94, 196, 265
- Green's theorem, 309
- Griewank's function, 173, 329

- Hamiltonian, 223
- hammer head, 29
- Hamming cliff, 91, 94
- Hamming distance, 94
- haploidy, 96
- hard bake, 7
- Helmholtz equation, 308
- heterozygosity, 96
- hierarchical fitness function, 179
- holographic lithography, 27
- homozygosity, 96
- Hopkins imaging, 211
- Hopkins' formula, 316
- hyperheuristic search, 151
- hyperplane, 91, 93, 110
- hypervolume indicator, 145

- i-line, 8
- ideal vector, 126
- illumination, 230
 - annular, 33
 - di-pole, 33
 - free-form, 33
 - Köhler, 17, 60, 230, 317
 - multi-pole, 33
 - off-axis, 30, 32
- image formation, 59
 - Abbe's approach, 65, 195, 210, 213, 230
 - Hopkins' approach, 20, 68

- vector imaging, 61, 73
- image log slope, 21, 211
 - normalized, 21
- image log-slope, 256
- image synthesis problem, 206
- imaging equation, 318
- immersion lithography, 21
- incidence angle-specific spectra, 252, 272
- incommensurability, 145
- incremental learning, 185
- indicator-based multi-objective evolutionary algorithms, 147
- integrated circuits, 1, 5
- integration function, 181
- intensification, 80
- intensity in the preferred polarization state, 265
- intensity peaks, 210
- interactive methods, 129
- interference mapping lithography, 210
- interference multilayer mirror, *see* multilayer
- intermediate focus, 22
- intermediate recombination, 82
- international technology roadmap for semiconductors, 1
- inverse lithography technique, 205, 223
- inverse problem, 205, 232
 - ill-posed, 205, 232
- ion-beam projection lithography, 26
- island models, 200
- iso/dense bias, 28
- joint eigenvector formulation, 209
- Jones calculus, 62
- Jones parameters, 265
- k-armed bandit, 119
- k-nearest neighbor, 142
- Karush-Kuhn-Tucker necessary condition, 163
- Kirchhoff approximation, *see* thin-mask approximation
- Kirchhoff mask approach, 229
- KKT condition, *see* Karush-Kuhn-Tucker necessary condition
- Krypton fluoride laser, 2
 - excimer, 8
- Lagrangian multipliers, 162
- Lamarckian evolution, 92, 150
- Lamarckian learning, 149
- Lambert's law, 74
- laser
 - Argon fluoride, 8
 - excimer, 8
 - exciplex, 8
 - Krypton fluoride, 8
 - molecular fluorine, 8
- level-set method, 77, 223
- Levenberg-Marquard method, 186
- lexicographical goals programming, 130
- LICQ, *see* linear independence constraint qualification
- light-beam search, 131
- line search, 165, 213
- line shortening, 29
- line-edge roughness, 26, 78
- line-end rounding, 29
- line-width roughness, 78
- line/space pattern, 28, 238
- linear classifiers, 180
- linear independence constraint qualification, 162
- linear programming technique, 218
- linear separability, 180
- linkage learning genetic algorithm, 122
- linkage-aware genetic algorithms, 122
 - linkage-learning genetic algorithm, 122
 - messy genetic algorithms, 122, 137
- LisBON is a Brokered Optimization Network, 202
- litho-friendly, 3
- load balance, 201
- local convergence, 165
- local search, 149
- local search candidate selection, 159
 - egalitarian, 159
 - elitarian, 159
 - tripartite, 159
- local search space restriction, 160
- local search trigger, 157
 - adaptive, 158
 - asymptotic, 158
 - dynamic, 158
 - exponential, 158
 - linear, 157
 - self-adaptive, 159
 - static, 157
- locus, 122
- Mack's model, 76
 - enhanced, 77
- Mansuripur transformation, 61
- manufacturability, 206, 207, 255, 288
- manufacturability criterion, 232
- marginal rate of substitution, 131
- marginal utility, 131
- Markov models, 124

- mask, *see* photomask
- mask absorber, 263
- mask absorber stack, 278
- mask diffraction variables, 209
- mask error enhancement factor, 18, 28, 205, 265
- mask-induced wavefront aberrations, 251
- mask-topography effects, 227, 263
- mating restriction, 112
- mating selection, 101
- Maxwell's equations, 55, 307
- mean-squared-error, 189
- meme, 149
- memetic algorithm, 149
- merit function, 152
- message passing interface, 200
- meta-Lamarckian learning, 152
- meta-search, 110
- metaheuristic, 79
- method of moments, *see* boundary element method
- model management, 179
- modulation transfer function, 20
- Moore's law, 1
- morphological regularization technique, 235
- mouse-bite, *see* anti-serif
- MPI, *see* message passing interface
- mse, *see* mean-squared-error
- multi-loss map, 128
- multi-meme memetic algorithm, 151
- multi-objective evolutionary algorithms, 133–147
 - performance assessment, 145
- multi-objective genetic algorithm, 135, 248, 256, 259
- multi-objective messy genetic algorithm, 137
- multi-objective optimization, 97, 101, 125, 198, 228, 238
- multi-objective struggle genetic algorithm, 135
- multilayer, 24
 - capping layer, 25
 - defects, 25
 - etch-stop layer, 25
- multilayer absorber, 264
- multilayer perceptron, 178, 179
- multimodal, 205, 210
- multimodality, 82
- multipole method, 56
- mutation, 79, 81, 89, 109, 120
 - point mutation, 109
- mutation rate, 110, 239
- mutual degree of coherence, 314
- mutual intensity, 314
- mutual intensity matrix, 225
- nanoimprint lithography, 27
- Nelder-Mead method, 80, 167, 174, 266
 - fortified descent, 171
 - restart methodology, 171
- Newton method, 162
- niched-Pareto genetic algorithm, 135
- niching, 112, 135, 139
 - clearing, 115
 - clustering, 116, 136, 143
 - crowding, 116, 135, 139
 - crowding distance, 139
 - deterministic crowding, 116
 - probabilistic crowding, 116
 - restricted tournament selection, 117
 - crowding factor groups, 118
 - multi-niche genetic crowding, 118
 - niching rule, 116
 - sequential niching, 113
 - sharing, 113, 135, 139
 - restricted tournament selection, 135
- NLP, *see* non-linear problem
- no-free-lunch theorem, 83, 96, 112, 122, 146, 149
 - free leftovers, 146
- no-preference methods, 129
- non-dominated ranking, 135
- non-dominated sorting genetic algorithm, 137, 138, 198
- non-linear problem, 161
- non-linear program, 213
- normal boundary intersection method, 132
- normal constraint boundary method, 132
- normalized image log-slope, 280
- notch model, 77
 - enhanced, 77
- numerical aperture, 8, 12
- objective function, 152
- objective parameter, 80
- objective space, 128
- objectives, 232
- Ohm's law, 307
- one-dimensional mask representation, 229, 233, 238
- one-point crossover, 196
- OPC, *see* optical proximity correction
- opening operator, 235
- optical proximity correction, 28, 206, 223
 - model-based, 30
 - rule-based, 30

- optimization parameter representation, *see*
decision space representation
- OptProb, 198
- outperformance, 145
- parameter control
 - adaptive, 110
 - deterministic, 110
 - dynamic, 110
 - self-adaptive, 110
- parameter representation, 229
- parametric source representation, 230
- paraxial approximation, 15
- Pareto dominance, 126
- Pareto efficiency, 126
- Pareto epsilon dominance, 126, 127
- Pareto front, 126, 128, 249, 296
- Pareto optimality, 126
- Pareto weak dominance, 127
- Pareto weak epsilon dominance, 127
- Pareto-archived evolution strategy, 136
- Pareto-envelope-based selection algorithm, 136
- Pareto-optimal set, 127
 - globally, 128
 - locally, 128
- partial differential equation, 75
- pattern doubling, 50
- pattern multiplication, *see* pattern doubling
- pay-off, 89
- penalty method, 222, 232
- penalty methods, 98
- perceptron, 179
- phase shifting mask, 229
- phase variations, 251
- phase-shifting mask, *see* phase-shifting mask, 206
- phenotype, 90, 149, 228
- photo-acid generator, 7, 74
- photo-active compound, 7, 74
- photomask, 5, 229, 263
 - alternating aperture phase-shifting mask, 44
 - alternating phase shifting mask, 278
 - alternating phase-shifting mask, 38
 - alternating phase-shifting masks, 38
 - attenuated phase shifting mask, 263, 278
 - high transmission, 242
 - low transmission, 242
 - attenuated phase-shifting mask, 38, 43
 - binary, *see* photomask, chromium on glass
 - bright field, 41
 - chromeless phase lithography, 38, 46
 - chromium on glass, 38
 - embedded attenuated phase-shifting mask, 44
 - phase conflict, 41
 - phase-shifting mask, 33, 38
 - rim phase-shifting mask, 42
 - strong phase-shifting mask, *see* photomask, alternating phase-shifting mask
 - topography, 55
 - trim mask, 41
 - weak phase-shifting mask, *see* photomask, attenuated phase-shifting mask
- photoresist, 5, 73, 74
 - chemically amplified, 7, 74
 - e-beam, 26
 - modeling, 73
 - acid dose diffusion threshold model, 78
 - diffused aerial image model, 78
 - lumped resist model, 78
 - mesoscopic models, 78
 - simplified models, 78
 - stochastic models, *see* mesoscopic models
 - threshold model, 78
 - negative tone, 5
 - novolak-DNQ, 5
 - positive tone, 5
 - two-photon, 51
- pixel-flipping technique, 206, 216
- pixelated mask, 209, 210, 229
- pixelated source, 210
- plasma
 - discharge-produced, 22
 - laser-produced, 22
- point-spread function, 225
- Poincaré sphere, 63
- polar-coordinate source representation, 231, 246
- polarization, 33, 61–63, 66, 73
 - p-polarized, *see* polarization, transverse magnetic
 - s-polarized, *see* polarization, transverse electric
 - transverse electric, 33, 61, 319
 - transverse magnetic, 33, 61, 319
- polyploidy, 96
- population, 89
 - size, 110
- population ranking, 138
- post optimization, 248
- post-bake, *see* hard bake
- post-exposure bake, 7, 75
- pre-bake, 5

- predator/prey, 99
- premature convergence, 222
- Price's theorem, 123
- process window, 16, 255, 263, 278
 - assymetry, 251
- projection lithography, 8
- projector wavefront, 251
- proximity curve, 241
- proximity lithography, 8
- pseudo-electric field, 226
- pseudo-meiosis, 96
- Python, 195
- Python multi-objective evolutionary algorithm, 198

- QP, *see* quadratic program
- quadratic program, 161, 213
- quadrupole illumination, 272
- quality indicators, 145
- quasi-Newton method, 162, 215, 227
- QuickProp, 186

- radial basis neural network, 178
- randomized weighted-sum-based genetic
 - algorithm, 134
- rank-based selection, 196
- Rastrigin's function, 172, 327
- RCWA, *see* rigorous coupled wave analysis
- recombination, 79, 89, 107, 120, 197
 - n-point crossover, 107
 - uniform crossover, 109
- recurrent neural network, 179
- reduced gradient method, 216
- reduced Hessian approximate, 165
- reference point methods, 131
- refractive index modulation, 293
- region-based selection, 137
- regularization, 218, 232
- reinforcement learning, 182
- repair method, 99, 232
- replacement scheme, 101
 - $(1 + 1)$, 81
 - $(1 + \lambda)$, 81
 - $(1, \lambda)$, 81
 - $(\mu + \lambda)$, 82
 - (μ, λ) , 82
- resilient backpropagation, 185, 266
- resist, *see* photoresist
- resist coating, 5
- resolution, 12–14
 - Rayleigh's criterion, 12–14
- resolution enhancement techniques, 27, 205
- response surface method, 178
- restricted tournament selection, 195, 239, 265
- RET, *see* resolution enhancement techniques
- rigorous coupled wave analysis, 55, 57
- Rosenbrock's function, 175, 325
- roulette wheel selection, 196, 239

- scalar imaging, 313
- scalarization, 125
- scaled defocus method, 71
- scanning electron microscope, 26
- scatter bar, *see* sub-resolution assist feature
- schema
 - defining length, 119
 - length, 119
 - order, 119
- schema growth, 120
- schema theorem, 91, 119
 - exact, 123
- schemata, 93
- Schwefel's function, 192, 327
- search space
 - continuous, 79
 - convex, 79
- selection, 79, 89, 101, 119
 - Boltzmann tournament selection, 106
 - Deterministic universal sampling, 104
 - rank-based, 103
 - remainder stochastic sampling, 104
 - roulette wheel, 101
 - stochastic sampling with replacement, 101
 - stochastic universal sampling, 103
 - tournament selection, 104
 - truncation selection, 106
- sensitivity, 288
- sequential quadratic programming, 161, 172
- serif, 29
- set coverage indicator, 145
- Shekel's foxholes function, 91, 172, 325
- side-wall angle, 16
- sidelobe criterion, 255
- sidelobe printing, 232, 246, 265
- sigmoid function, 218
- simplex, 167
 - contraction, 167
 - expansion, 167
 - greedy, 168
 - inside contraction, 169
 - outside contraction, 169
 - reflection, 167
 - shrink, 167
- simulated annealing, 79
- single-layer perceptron, 180
- singular value decomposition, 68, 70

- size bias optical proximity correction, 254
- slope criterion, 237
- SMO, *see* source/mask optimization
- Sobolev norm, 211
- source
 - extended, 16
 - finite, 8
- source/mask optimization, 205, 207
- SQP, *see* sequential quadratic programming
- statistical mechanics models, 125
- steady state GA, 197
- steady-state genetic algorithm, 101
- steepest decent, 220
- step-and-repeat, 10
- step-and-scan, 10
- stochastic ranking, 99
- Stokes parameters, 63
- strategy parameter, 80
- stratified medium, 73
- strength Pareto evolutionary algorithm, 198
- strength pareto evolutionary algorithm, 140, 248
- strength-Pareto evolutionary algorithm, 137, 296
- strict Pareto dominance, 145
- strict partial order, 126, 127
- sub-resolution assist feature, 29
- sub-resolution assist features, 223
- sum of coherent systems, 69, 70, 206, 210, 218, 226
- supervised learning, 182
- surface inhibition, 76
- survivor selection, 101
- swarm intelligence, 88
 - ant colony optimization, 89
 - artificial bees colony optimization, 89
 - bees algorithm, 89
 - cuckoo search, 89
 - firefly algorithm, 89
 - particle swarm optimization, 88
- swing curve, 291
- tabu search, 79
- Tau-map, 225
- thin-film optics, 73
- thin-mask approximation, 55
- threshold intensity, 16
- threshold material, 51
- through-focus evaluation, 237
- through-focus optimization, 223
- topological mask effects, 251
- tournament selection, 196, 239
- tracks-and-sectors source representation, 231, 246
- transfer matrix, 321
- transfer matrix method, 71, 319
- transmission cross-coefficient, 20, 68
- traveling salesman problem, 96
- tri-objective optimization, 299
- triangular decomposition, 209
- trim exposure, *see* photomask, trim mask
- trust region approach, 166
- Twisted matrix, 201
- two-beam imaging, *see* two-beam interference
- two-beam interference, 14
- two-dimensional mask representation, 229, 230, 233, 241
- two-point crossover, 196, 265
- uniform crossover, 196
- uniform intensity distribution, 211
- unsupervised learning, 182
- utility function method, 130
- utopian vector, 126
- vector-evaluated genetic algorithm, 134
- Vose's model, *see* dynamic systems models
- wafer
 - stack, 71
 - topography, 71
- Walsh transform, 123
- wave equation, 308
- wavefront aberration, 10, 60
- wavefront control, 254
- wavefront engineering, 207
- waveguide method, 57, 195, 252, 272, 291
- wavelets, 222
 - Haar, 222
- weak Pareto dominance, 145
- weighted sum approach, 125, 131, 237, 265, 291
- XML-RPC, 202
- XOR gate, *see* exclusive-or gate
- Zernike polynomials, 60, 254