Shape Dithering for 3D Printing

MOSTAFA MORSY ABDELKADER MORSY, Fraunhofer Institute for Computer Graphic Research IGD, Germany and Technical University Darmstadt, Germany
ALAN BRUNTON, Fraunhofer Institute for Computer Graphic Research IGD, Germany
PHILIPP URBAN, Fraunhofer Institute for Computer Graphic Research IGD, Germany and Norwegian University of Science and Technology NTNU, Norway

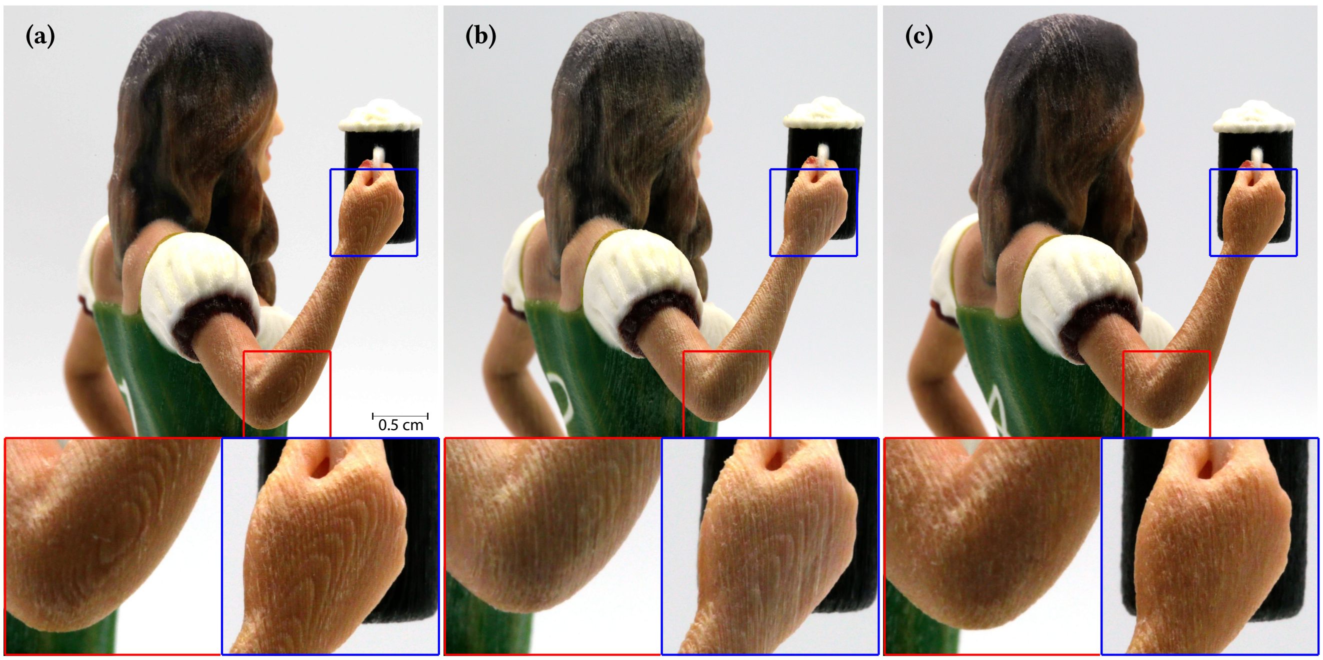

Fig. 1. Quantization artifacts are a fundamental issue for all 3D printing technologies, but especially for multi-material jetting printers. Despite their high resolutions, staircasing artifacts (a) can be visually irritating, and can structurally weaken the part [Moore and Williams 2015]. Existing techniques [Kritchman 2010] (b) are limited to specific surface orientation, introduce considerable extra computation, and do not remove all artifacts. Our purely geometric and algorithmic technique (c) removes staircase artifacts in all surface orientations, accounts for resolution anisotropy, and introduces a minimal computational overhead.

We present an efficient, purely geometric, algorithmic, and parameter free approach to improve surface quality and accuracy in voxel-controlled 3D printing by counteracting quantization artifacts. Such artifacts arise due to the discrete voxel sampling of the continuous shape used to control the 3D printer, and are characterized by low-frequency geometric patterns on surfaces of any orientation. They are visually disturbing, particularly on small prints or smooth surfaces, and adversely affect the fatigue behavior of printed parts. We use implicit shape dithering, displacing the part's signed distance field with a high-frequent signal whose amplitude is adapted to the (anisotropic) print resolution. We expand the reverse generalized Fourier slice theorem by shear transforms, which we leverage to optimize a 3D blue-noise mask to generate the anisotropic dither signal. As a point process it is efficient and does not adversely affect 3D halftoning. We evaluate our approach for efficiency, geometric accuracy and show its advantages over the state of the art.

Authors' addresses: Mostafa Morsy Abdelkader Morsy, Fraunhofer Institute for Computer Graphic Research IGD, Darmstadt, Germany and Technical University Darmstadt, Darmstadt, Germany, mostafa.morsy.abdelkader.morsy@igd.fraunhofer.de; Alan Brunton, Fraunhofer Institute for Computer Graphic Research IGD, Darmstadt, Germany, alan.brunton@igd.fraunhofer.de; Philipp Urban, Fraunhofer Institute for Computer Graphic Research IGD, Darmstadt, Germany and Norwegian University of Science and Technology NTNU, Trondheim, Norway, philipp.urban@igd.fraunhofer.de.

0730-0301/2022/7-ART82 $15.00
https://doi.org/10.1145/3528223.3530129

CCS Concepts: • **Computing methodologies** → **Shape modeling**; *Perception.*

Additional Key Words and Phrases: 3D printing

ACM Reference Format:
Mostafa Morsy Abdelkader Morsy, Alan Brunton, and Philipp Urban. 2022. Shape Dithering for 3D Printing. *ACM Trans. Graph.* 41, 4, Article 82 (July 2022), 12 pages. https://doi.org/10.1145/3528223.3530129

1 INTRODUCTION

We present efficient, purely geometric and algorithmic techniques to improve surface quality and accuracy by counteracting quantization artifacts in 3D printing.

Quantization is a fundamental aspect of digital fabrication, and additive or layered technologies in particular. Even technologies that are driven by vector paths have motor tolerances that prevent truly continuous precision, and all layer-based technologies require discrete slices of finite thickness [Alexa et al. 2017]. In this paper, we focus on multi-material jetting systems, which are controlled on a voxel basis, introducing an explicit quantization of the input geometry along all three axes.

Multi-material jetting (*a.k.a* polyjetting) is the most versatile 3D printing technology for creating spatially-varying mechanical and visual effects. Printing systems from Mimaki [2020] or Stratasys [2020] can combine 6 or 7 printing materials with different optical and mechanical properties into a single print. This allows creating spatially-varying subsurface scattering [Dong et al. 2010; Hašan et al. 2010], full-color [Babaei et al. 2017; Brunton et al. 2015] or joint color and translucency [Brunton et al. 2018]. Multi-material jetting becomes more affordable [3D Printing Industry 2020], faster [DP Polar 2020] and the printed parts significantly more durable. The latter is possible due to novel print-head technologies allowing jetting of high-viscose printing materials [Quantica 2020; Xaar 2020]. The achievable part durability expands the application areas of material jetting to manufacture parts that withstand high mechanical load in combination with desired spatially-varying appearance, such as dental crowns or implants. Approaches were proposed to compensate adverse effects caused by intrinsic optical printing material properties, such as light transport texture blurring [Elek et al. 2017; Sumin et al. 2019]. But only little work was published focusing on correcting or reducing artifacts caused by the binary way such printing systems are controlled or due to the printing process itself.

1.1 Quantization errors

The voxel resolution of binary 3D printers (such as material jetting) defines the upper limit of geometric accuracy since just the whole volume of the voxel can be filled with build material or left empty.

Resulting quantization errors between the input surface and the voxel surface may (and mostly) include low-frequency patterns and edges, i.e. the quantization signal on the surface contains directional high and low-frequency components. Since most 3D printing systems employ anisotropic voxels, the magnitude of quantization errors depends on the surface orientation.

The directional, surface-orientation-dependent quantization errors yield staircasing artifacts in the final 3D print. These artifacts are visually disturbing when viewed at close range in particular for small objects due to the low signal-to-noise-ratio. For larger objects viewed from a greater distance they can create unwanted specular highlights at surfaces viewed under off-specular conditions.

Section 5 shows quantitatively through simulation and visually with printed parts how staircasing contributes to surface roughness and our proposed dithering approach effectively compensates for this.

In addition to visual effects, staircasing artifacts may adversely affect the fatigue behavior of parts. Even though we are not evaluating fatigue behavior in this paper, studies have shown that lower surface roughness due to glossy as opposed to matte finishing improves tensile strength [Kampker et al. 2017] and fatigue behavior [Moore and Williams 2015]. Moore and Williams [2015] further suggested that staircasing (that is reduced by our dithering approach) introduces surface roughness that negatively contributes to fatigue behavior.

1.2 Voxel representation vs. real print

A voxel representation is an idealistic model of material placement. In addition to droplet-positioning noise and pre-cured material mixing, the printer's firmware places dots in interlaced patterns [Napadensky 2014] with considerable overlap between voxels resulting in material cross-contamination between neighboring voxels. This is necessary to ensure structural integrity of the print. The correct Z-height is enforced by leveling mechanisms such as a roller or a squeegee, which push excess material into neighboring voxels. Furthermore, build and support material mix at vertical surfaces before curing so that not all support material can be removed in the cleaning process [Brunton et al. 2015]. These printing-process effects are essentially a low-pass filter of the quantization signal but also introduce stripes orthogonal to the printer's Z-axis. To obtain pleasant surfaces, some previous works have surrounded the objects with a coating of clear material, then removed staircase artifacts and remaining support manually in a post-process by sanding or polishing [Elek et al. 2017; Sumin et al. 2019].

1.3 Contribution

In this paper, we propose an efficient, purely geometric, and parameter free approach to improve surface quality and accuracy, in particular to reduce quantization artifacts. We make the following technical contributions:

(1) An implicit formulation of a shape dithering algorithm as a point process using displacements, which supports streaming and distributes the quantization error over the part's surface reducing staircasing artifacts of the final print.
(2) The expansion of the reversed generalized Fourier slice theorem to shear transforms (and its proof) used to optimize an isotropic 3D blue-noise mask employed in 1 to be applicable to anisotropic printer resolutions.

Since our approach is completely geometric and algorithmic in nature, we can apply and tune it to any make and model of Polyjetting 3D printer, with no special hardware or firmware required.

2 RELATED WORK

2.1 Dithering

In binary printing dithering is used to create binary patterns determining ink or material distributions for reproducing full-tone input. Most algorithms are developed for the purpose of creating visually pleasing full-tone reproductions by patterns that shift quantization errors to high spatial frequencies for which the human visual system is less sensitive.

We can classify dither methods into one of three categories: *point processes*, *neighborhood processes* and *iterative processes*. The fastest and most efficient methods belong to the point process category. They use a precomputed dither mask tiled over the input signal to simply threshold the signal by the corresponding mask entry to create the binary pattern. Small masks were used for dispersed-dot ordered dither [Bayer 1973] with the drawback of cross-hatch pattern artifacts caused by the mask tiling. Cho *et al.* extended the approach to 3D printing [Cho et al. 2003]. To minimize artifacts, much bigger masks were developed and used, such as blue-noise masks [Mitsa and Parker 1992; Ulichney 1993] or green-noise masks [Lau et al. 1999]. Blue-Noise masks create quantization errors with minimal low-frequency content minimizing apparent artifacts. The quantization errors resulting from green-noise masks are dominated by mid-frequencies taking into account limited accuracy of dot placement by the printing system (e.g. laser printers). Morovič *et al.* used blue- and green-noise masks to create patterns for multi-material printing, which are extendable to 3D printing [2017a; 2017b].

A neighborhood processes uses the pixel or voxel's spatial neighborhood for thresholding the input signal. Floyd and Steinberg [1976] introduced error diffusion that creates high-frequency quantization errors: The method traverses the pixels of an image, thresholds the signal at a pixel with a value of 0.5 and adds the quantization error to its neighbors using an error diffusion filter. Error diffusion was further optimized to reduce remaining low-frequency structural artifacts [Chang et al. 2009; Ostromoukhov 2001; Zhou and Fang 2003]. Brunton *et al.* extended error diffusion to 3D printing [2015] by traversing the object's offset-voxel-surfaces, thresholding the signal at a voxel and distributing the quantization error to its neighbors using an oriented anisotropic error diffusion filter. Alexa and Kyprianidis have adapted error diffusion to operate on meshes [2015]. Using error diffusion on meshes to displace the surface would require refining the tessellation so that triangles are on the order of the printer's native resolution before applying error diffusion.

An iterative process creates an initial binary pattern and iteratively modifies it to minimize the distance to the full-tone input. This distance can be computed via image difference metrics as used in the direct binary search approach [Agar and Allebach 2005; Pang et al. 2008]. Iterative processes produce high visual quality patterns, but are in theory and practice the slowest of the three categories.

2.2 Surface manipulation for 3D printing

Luongo *et al.* modulated the microstructure of the print's surface to control its roughness using a smooth noise function changing the surface's voxel filling rate [2020]. This is possible for printing technologies such as Digital Light Processing (DLP) or Liquid Crystal Display (LCD) stereolithography, allowing to control the amount of light exposed onto voxels via grayscale values. Luongo *et al.* used anti-aliasing to remove staircasing artifacts and could create spatially-varying specular reflectance. Anti-aliasing increases geometric resolution and therefore reduces quantization errors by modulating the voxel's filling-rate. This is not possible for printing systems that can only be controlled by binary patterns allowing to completely fill a voxel or to leave it empty. In contrast our dithering approach distributes the quantization error spatially and shifts it to higher spatial frequencies that are partly removed by the low-pass characteristics of the printing process. Orth *et al.* corrected for ray distortion in tomographic 3D printing by re-sampling the parallel-beam radon transform into an aberrated geometry [2021]. This technique is specific to this printing technology.

Page *et al.* [2017] modeled the height modulation of relief 3D prints by convoluting the digital input by a Height Modulation Transfer Function (HMTF) measured using a white light confocal microscope. They corrected the digital input to obtain the intended prints by deconvolution via a Wiener filter.

Alexa *et al.* [2017] proposed a technique for optimizing the slice thickness, for technologies that allow to vary the slice thickness during printing, primarily stereolithography (SLA) and fused filament fabrication (FFF). This technique finds a global optimum of slice thicknesses in terms of the volumetric error of a voxel representation of the input surface. Polyjetting technologies currently do not support varying the slice thickness within a single print.

Kritchman [2010] proposed many techniques related to improving surface quality and geometric accuracy in material-jetting 3D printing, in particular to prevent print deformation, but also a technique to reduce quantization artifacts by slicing at double the native resolution along the axis for which the printer resolution is the lowest, and using an interlacing approach to reduce the output back to the native resolution. This is a directional geometric dither that is effective in addressing the voxel anisotropy. However, its effectiveness is limited to distinct surface orientations, and it is computationally expensive since the number of addressed voxels doubles. To our knowledge, this is the only previous work that addresses quantization errors in material-jetting 3D printing. It is therefore our primary point of comparison in Section 5, where we also detail how interlacing can be understood as a directional dither process.

Various methods for surface manipulation focus on creating or reproducing distinct optical effects in printing. Baar *et al.* [2014] changed the deposition time between printing two layers and Samadzadegan *et al.* [2015] as well as Piovarči *et al.* [2020] used varnish halftones to create specially varying gloss. Malzbender *et al.* printed on a structured surface and used dispersed dithering to selectively place ink for creating surface reflectance [2012]. Matusik *et al.* [2009] used error diffusion to place absorption and metallic inks on a substrate to produce spatially-varying bidirectional reflectances. Lan *et al.* [2013] used a 3D printer to create oriented facets and a color 2D printer to color these facets for reproducing spatially-varying bidirectional reflectance incl. anisotropic reflectance. Rouiller *et al.* [2013] used also two printers to fabricate small domes with a polyjetting printer and attach them to a geometry produced by a binder jetting printer. By varying the domes' geometry bidirectional reflectance can be adjusted locally. The approaches of Lan *et al.* or Rouiller *et al.* are limited to shapes without self occlusions with a relative small curvature. The microfacet or dome resolution must be significantly lower than the print resolution to allow for shape variations of facets/domes, i.e. the approaches cannot be used to reduce the visibility of quantization errors.

Auzinger *et al.* proposed an optimization framework to create structural colorizations of surfaces by multi-photon lithography employing finite-difference time-domain simulations [2018]. Since controlling material solidification in multi-photon lithography is not restricted to a fixed voxel grid, this technology does not show quantization errors as addressed in this paper.

3 IMPLICIT SHAPE DITHERING

As described in Section 2.1, a dithering process introduces a secondary signal to cancel or distribute quantization errors. The dither signal is typically a high-frequency signal, *e.g.* a blue-noise signal, as this distributes the quantization errors to high frequencies, which are removed by the low-pass filtering of the human visual system, and in fact many 3D printing processes. The quantization process for binary 3D printers is voxelization: conversion of the conceptually continuous shape to a discrete voxel grid. The quantization error is the difference between these two representations. For background and nomenclature concerning voxelization we refer the reader to Cohen-Or and Kaufmann [1995].

Let $\mathcal{S} \subset \mathbb{R}^3$ be a closed set representing a shape and $\partial\mathcal{S} \subset \mathcal{S}$ its surface. Let the build space of the printer $\mathcal{B} \subset \mathbb{R}^3$ be a semi-open cuboid volume decomposed with an axis-aligned, fixed, regular grid of disjoint voxels C at printer resolution so that $\bigcup_{v \in C} v = \mathcal{B}$ and $u \cap v = \emptyset$ for $u, v \in C, u \neq v$. In practice, we use $\mathcal{B}(\mathcal{S}) \subset C$, a padded voxel-bounding-box of $\mathcal{S}$ s.t. $\mathcal{S} \subset \mathcal{B}(\mathcal{S})$. Padding need only be on the order of a few voxels, as described below.

For binary 3D printers the control signal sent to the printer is represented by a material assignment $\mathbf{P}(v) : \mathcal{B}(\mathcal{S}) \mapsto \mathcal{M}$, where $\mathcal{M}$ is a set of possible material assignments, including the assignment of no material (empty) to a voxel. The set of voxels $\mathcal{V} \subset \mathcal{B}(\mathcal{S})$ for which we assign a non-empty material, $\mathcal{V} = \{v \in \mathcal{B}(\mathcal{S}) : \mathbf{P}(v) \neq \text{empty}\}$, is our quantized approximation of $\mathcal{S}$ to be printed.

Directly computing $\mathcal{V}$ by a solid voxelization of $\mathcal{S}$ results in a quantization error with spatially coherent steps, staircasing artifacts, which are rounded, but still spatially coherent following the low-pass effect of the printing process. This creates visual and structural artifacts as described in Section 1.1.

To create a voxel representation $\mathcal{V}$ of $\mathcal{S}$, which contains high-frequency surface topography that under the printing process' low-pass effect results in a smooth surface without coherent steps, we can use dithering. To dither we need a signal, and we could start by computing a tonal function $t : \mathcal{B}(\mathcal{S}) \mapsto \mathbb{R}$, where $t(v) = 1$ for voxels contained entirely inside $\mathcal{S}$, $t(v) = 0$ for voxels entirely outside $\mathcal{S}$, and $0 < t(v) < 1$ for voxels intersected by $\partial\mathcal{S}$.

Next, we could dither or halftone t to get a binary segmentation $\mathcal{V}$. An established technique for this is to use a dither mask $\mathbf{M} : \mathbb{R}^3 \mapsto \mathbb{R}$ to threshold t

$$v \in \mathcal{V} \iff t(v) > \mathbf{M}(v) \tag{1}$$

where a careful choice of $\mathbf{M}$ gives the high-frequency structure of $\mathcal{V}$ that results in the desired surface under the printing process.

Instead of using a tonal function, we consider the scenario where we are given as input the surface $\partial\mathcal{S}$ rather than the solid shape, as this represents the standard workflow for 3D printing. We convert the input representation of $\partial\mathcal{S}$, *i.e.* a triangle mesh, to an implicit representation, $\mathbf{d} : \mathcal{B}(\mathcal{S}) \mapsto \mathbb{R}$. We can then determine the non-empty voxels $\mathcal{V}$ by thresholding: $\mathcal{V} = \{v \in \mathcal{B}(\mathcal{S}) : \mathbf{d}(v) < \tau\}$, for some threshold τ. A popular and intuitive choice for $\mathbf{d}$ is the signed distance to $\partial\mathcal{S}$,

$$\mathbf{d}(v) \quad = \quad \mathbf{s}(v) \min_{x \in \partial\mathcal{S}} \|x - \mathbf{c}(v)\|_2 \tag{2}$$

where $\mathbf{c}(v)$ is the centroid of voxel v and $\mathbf{s}(v) = -1$ if $\mathbf{c}(v) \in \mathcal{S}$ and otherwise $\mathbf{s}(v) = 1$.

Directly thresholding $\mathbf{d}$ with $\tau = 0$ will result in the same quantization errors as directly computing a solid voxelization (modulo an erosion by up to a voxel). We could compute a tonal function t from $\mathbf{d}$ and threshold this with $\mathbf{M}$, as described above, but this would require computing the tonal function and reading the mask at all voxels of $\mathcal{B}(\mathcal{S})$. Since we are interested in adding a high-frequency topography to the surface, an equivalent approach is to use $\mathbf{M}$ to generate a high-frequency signal with which to perturb the surface and threshold with a fixed value. Displaced signed distance fields [Brunton and Abu Rmaileh 2021] provide an efficient and accurate approximation of the true signed distance fields, which further allow to introduce a spatially varying displacement, or offset, thereby controlling the 0-level set of the implicit function. Using this framework, we can *implicitly* dither the surface by adding a spatially varying offset $\mathbf{f}$ to $\mathbf{d}$,

$$\mathbf{d}'(v) \quad = \quad \mathbf{d}(v) + \mathbf{f}(\bar{v}) \tag{3}$$

where

$$\bar{v} \quad = \quad \underset{u \in \partial\mathcal{V}}{\operatorname{argmin}} \|\mathbf{c}(v) - \mathbf{c}(u)\|_2 \,. \tag{4}$$

Thresholding $\mathbf{d}'$ with $\tau = 0$ gives our dithered shape $\mathcal{V}'$.

By carefully defining and computing $\mathbf{f}$ we can distribute the quantization error of $\mathcal{V}'$ uniformly over the surface with high spatial frequencies. We define $\mathbf{f} : \partial\mathcal{V} \mapsto \mathbb{R}$ as

$$\mathbf{f}(v) \quad = \quad 4\, \mathbf{k_D}(v)\, (\mathbf{M}(v) - 0.5) \tag{5}$$

where $\mathbf{k_D} : \partial\mathcal{V} \mapsto \mathbb{R}^+$ is a function adjusting the displacement magnitude according to the surface orientation and the anisotropic voxel dimensions d_x, d_y, d_z, and $\mathbf{M} : \partial\mathcal{V} \mapsto [0, 1]$ is a high-frequency noise function on the surface voxels and the constant 4 scales the signal's maximum amplitude to the voxel diameter along the surface normal. The latter applies for

$$\mathbf{k_D}(v) \quad = \quad \frac{1}{2\,\|\mathbf{Dn}(v)\|_\infty} \tag{6}$$

where $\mathbf{n}(v) = (n_x, n_y, n_z)^T$ is the surface normal for voxel v and $\mathbf{D} = \operatorname{diag}(1/d_x, 1/d_y, 1/d_z)$. The function $\mathbf{k_D}$ returns the distance between voxel center $\mathbf{c}(v)$ and boundary of the voxel along the direction $\mathbf{n}(v)$.

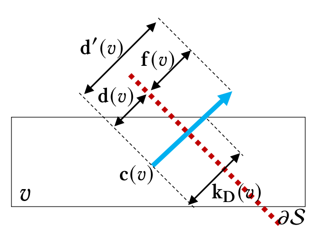

The number of voxels for padding $\mathcal{B}(\mathcal{S})$ is 1. Note that displacing the surface according to (3) is equivalent to creating the tonal function

$$t(v) = \max(-1, \min(1, -\mathbf{d}(v)/4\mathbf{k_D}(v))) + 0.5 \tag{7}$$

and thresholding with $\mathbf{M}$ as in (1).

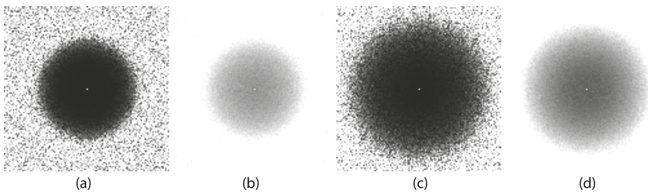

Fig. 2. Slice through the 3D Fourier transform of a 128^3 blue noise mask computed via the void and cluster method [Ulichney 1993] using code by [Peters 2017] employing (a) std = 1.9 and (c) std = 1.1. Axis-aligned 1D integral projection of the 3D Fourier transform: (b) of the mask (a), (d) of the mask (c). Absolute FFT values normalized by the mean and clipped at 1.

4 APPROXIMATING BLUE NOISE ON SURFACES

To endow $\mathbf{f}$ with the highest possible spatial frequencies on $\partial\mathcal{V}$, we would ideally like $\mathbf{M}$ to have blue noise characteristics w.r.t. the metric of $\partial\mathcal{S}$. That is, the projection of $\mathbf{M}$ onto the eigenfunctions of the Laplace-Beltrami operator of $\partial\mathcal{S}$ would result in minimal low-frequency components. While works exist to perform blue noise *sampling* on surfaces [Yan et al. 2015], to our knowledge, generating blue noise functions on surfaces remains a challenging open problem. In particular, even techniques for blue noise sampling on surfaces do not approach the efficiency of point processes for dithering. We experimented extensively with a technique to exploit the **Generalized Fourier Slice Theorem** introduced and proven by Ng [Ng 2005] by integral projections of 3D blue noise masks. However, we found this technique to offer insufficient improvement in dither quality for the associated computational cost, when compared to direct use of optimized 3D blue noise masks as described below. The interested reader is referred to the supplemental material for more details on the integral projection method.

4.1 Limitations of a 3D Blue Noise (BN) Mask to produce 2D Blue Noise on Surfaces

Figure 2(a) shows a slice of a 3D BN mask's Fourier transform, illustrating the negligible low-frequency content within a sphere centered at the DC coefficient (dark disk) and large high-frequency content at the borders. According to Corollary 1 (Reverse Generalized Fourier Slice Corollary), the Fourier transform of the signal sampled from a 3D BN mask on an arbitrarily oriented plane is similar to an integral projection of the BN mask's Fourier transform onto the plane. As can be seen from Figure 2(a) an integral projection in Fourier domain will always add high-frequency content from the borders to the low-frequencies in the center resulting in more low-frequency content. Therefore, sampling a 3D BN mask at an object's surface results in a distribution with reduced blue-noise character. This is illustrated in Figure 2(b), where the low-frequency content of the signal sampled from a 3D BN mask on an axis-aligned plane is considerable larger than in (a). Peters showed this for an axis aligned slice of a 3D BN mask [Peters 2017], Lagae and Drettakis introduced and proved the slice-projection theorem for stochastic processes [2011] that generalized it to any orientation and dimension. We expand this theorem by shear transforms according to Corollary 1, which allows us to optimize an isotropic 3D blue-noise mask to be applicable to anisotropic printer resolutions since shearing a space is equivalent to rotating and dilating the space.

Corollary 1. ***(Reverse Generalized Fourier Slice Corollary):*** *Let B be an invertible NxN-dim. matrix defining a rotation and/or shear transform, then the M-dim. Fourier transform F^M of the M-dim. slice S_M^N oriented according to B through the N-dim. function g is equal to the M-dim. integral projection I_M^N onto a similarly oriented plane of the Fourier transform of g, i.e. (see Fig. 3)*

$$\left(F^M \circ S_M^N \circ B\right)[g] = \left(I_M^N \circ \left(B^{-T}/|B^{-T}|\right) \circ F^N\right)[g] \tag{8}$$

The exact definition of the symbols and the proof of the corollary is given in Appendix A. In our case $N = 3$ and $M = 2$.

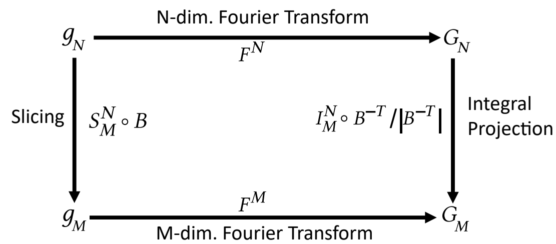

Fig. 3. Reverse Generalized Fourier Slice Corollary: Transform relationships between an N-dim. function g_N, their M-dim. slicing g_M, and the subsequent M-dim. Fourier transform.

4.2 Optimizing the 3D Blue Noise Mask

In order to minimize low-frequency content of 2D signals sampled from a 3D mask at a plane of any orientation, we tune (scale) the mask's cutoff frequency [Mitsa and Parker 1992]. This is possible by using the standard deviation in Ulichney's void and cluster algorithm [Ulichney 1993]. Ulichney recommended a standard deviation of 1.5 for a 2D BN mask and Peters used 1.9 [Peters 2017]. The smaller the standard deviation the higher is the scaled cutoff frequency and the integral projection results in smaller low-frequency content. However, a too high cutoff frequency creates low-frequency leakage as observed by Mitsa and Parker [Mitsa and Parker 1992], which increases low-frequency content. We compute a set of 32^3 BN masks M_σ with standard deviations $\sigma = \{0.5, 0.6, \ldots, 3.1\}$ from which we select an optimal mask for our purpose. For this, we compute the *Radially Averaged Power Spectral Density* (RASPD) [Lau and Arce 2001] of the 2D signal resulting from slicing mask M_σ along a surface oriented according to B_i employing Corollary 1:

$$\mathrm{RASPD}_{\sigma,i} = \mathrm{RASPD}\left(\left|\left(I_3^2 \circ \left(B_i^{-T}/|B_i^{-T}|\right) \circ F^3\right)[M_\sigma]\right|\right) \tag{9}$$

$\mathrm{RASPD}_{\sigma,i}$ is a vector with elements storing the radially averaged densities corresponding to a set of annular ring frequency bands [Lau and Arce 2001]. We use 32 bands in this study. We select a set of B_i to cover all surface orientations defined by normals sampling the unit sphere in 1° polar and azimuthal angle intervals. To obtain a single RASP reflecting the worst low-frequency content of all such slices through the 3D mask M_σ from $\mathrm{RASPD}_{\sigma,i}$ considering potential anisotropies we compute

$$\overline{\mathrm{RASPD}}_\sigma(f) = \max_i \mathrm{RASPD}_{\sigma,i}(f) \tag{10}$$

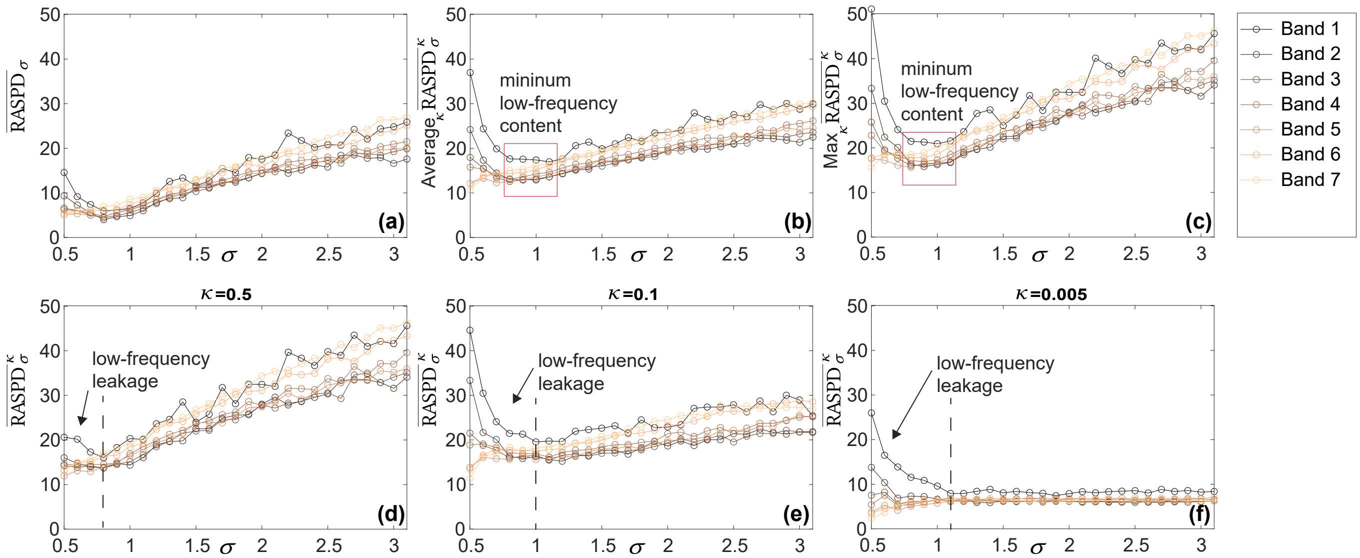

Fig. 4. $\overline{\mathrm{RASPD}}_\sigma$ and $\overline{\mathrm{RASPD}}^\kappa_\sigma$ plots for the lowest frequency bands (band i contains smaller frequencies than band i+1). The dashed lines in (d)-(f) indicate the σ-threshold at which low-frequency leakage affects the masks of smaller σ-values.

where $\mathrm{RASPD}_{\sigma,i}(f)$ is the RASPD value at frequency band f. The first seven smallest frequency bands are shown in Fig. 4(a). We see that for high σ values low-frequency content is high. The mask with $\sigma = 0.8$ shows the minimum low-frequency content.

In our application the surface signal is quantized by the voxel grid. Dithering the surface translates in thresholding M_σ at a distinct level $\kappa \in [0, 1]$. Therefore, we investigate the RASPD of the thresholded masks M^κ_σ, which are the binary masks resulting from thresholding M_σ, i.e. $M^\kappa_\sigma \doteq 1$ for $M_\sigma > \kappa$ and $M^\kappa_\sigma \doteq 0$ else, where $\doteq$ is the element-wise equality relationship. The objective for dithering is to find a mask M_σ containing minimum low-frequency content when thresholded by any $\kappa \in [0, 1]$. This is because low frequency content introduces voids and clusters of voxels which cannot be removed by the low-pass filter of the printing process. This results in unwanted meso surface-roughness. We compute $\overline{\mathrm{RASPD}}^\kappa_\sigma$ by inserting M^κ_σ in Eq. (9) and (10). Fig. 4(b) and (c) show low-frequency bands of averaged and maximized $\overline{\mathrm{RASPD}}^\kappa_\sigma$ w.r.t. κ. Masks corresponding to $\sigma = 0.8, 0.9, \ldots, 1.1$ contain the smallest low-frequency content and perform almost similarly w.r.t. RASPD.

We observed that the more κ deviates from 0.5 the greater the low-frequency leakage of masks of larger σ as shown in Fig. 4(d-f). Fig. 5 shows how this impacts voids and clusters within slices through thresholded masks. We selected $M_{\sigma=1.1}$ as the optimal mask for our purpose since it is not affected by low-frequency leakage for all κ values and is near the $\overline{\mathrm{RASPD}}^\kappa_\sigma$ minimum.

5 EXPERIMENTS

We tested our approach on a Mimaki 3DUJ-553 material-jetting printer using the standard cyan, magenta, yellow, black, white and clear inks. The white and black printing materials have high opacity making staircasing artifacts particularly apparent. The voxels are anisotropic with $(d_x, d_y, d_z) \approx (42, 84, 22)$ μm. We created prints

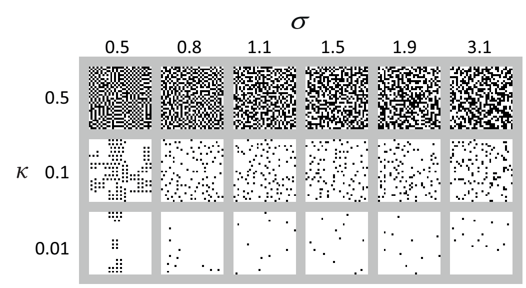

Fig. 5. Slices through masks M_σ thresholded by κ

by combining our techniques with the implementation provided by the Cuttlefish SDK [IGD 2020] for color gamut mapping, separation, halftoning, and printer-specific output. We evaluated our method with 5 models, shown in Figure 6. The three in the top row were printed to be 12cm along their longest axis. The longest axis was aligned with the X-axis of the printer. The cube was rotated by 2° about all principle axes. We carefully cleaned all prints by dissolving the support in water. No polishing, sanding or coating were applied.

To validate the importance of our algorithmic choices, we compare our approach against direct voxelization with no staircasing correction (“Control"), interlacing [Kritchman 2010], implicit dithering with white noise and implicit dithering with a blue-noise mask with a sub-optimal (according to the measures used in Section 4.2) standard deviation of 2.5. In the supplemental material we discuss why common procedural noise functions as well as error diffusion methods are not well suited for creating the dither signal.

Interlacing is the most relevant existing work: To our knowledge, it is the only existing work addressing quantization in material-jetting 3D printers, and it can be understood as a dithering process. In this formulation, the mask/threshold is constant, *e.g.* 0.5, and a

Spirit Rider
$|\mathcal{B}| \approx 6.9 \cdot 10^9$
$|\mathcal{V}| \approx 0.66 \cdot 10^9$
$|\partial\mathcal{V}| \approx 12.2 \cdot 10^6$

Queen
$|\mathcal{B}| \approx 4.0 \cdot 10^9$
$|\mathcal{V}| \approx 0.81 \cdot 10^9$
$|\partial\mathcal{V}| \approx 7.4 \cdot 10^6$

Dirndl
$|\mathcal{B}| \approx 2.3 \cdot 10^9$
$|\mathcal{V}| \approx 0.31 \cdot 10^9$
$|\partial\mathcal{V}| \approx 6.7 \cdot 10^6$

Sphere
$|\mathcal{B}| \approx 1.0 \cdot 10^9$
$|\mathcal{V}| \approx 0.43 \cdot 10^9$
$|\partial\mathcal{V}| \approx 3.1 \cdot 10^6$

Cube
$|\mathcal{B}| \approx 1.2 \cdot 10^9$
$|\mathcal{V}| \approx 0.81 \cdot 10^9$
$|\partial\mathcal{V}| \approx 6.0 \cdot 10^6$

Fig. 6. Prints of our test models and their approximate voxel counts.

possible tonal value is

$$t(v) = \begin{cases} 1 & \text{if } \mathbf{d}(v + [0\ 0.25d_y((s \mod 2) - 1)\ 0]^T) < 0 \\ 0 & \text{otherwise} \end{cases} \tag{11}$$

where s is the integer slice number, and we assume sampling at voxel centroids. Note that t is sampled on a voxel grid with spacing (d_x, d_y, d_z) whereas the signed distance $\mathbf{d}$ is sampled on a voxel grid with double resolution in y, hence the $\pm$ quarter voxel offset in y for alternating slices. Comparisons to white noise and sub-optimal blue noise were included to show the importance of a high-frequency dither signal and, in the case of white noise, to put the efficiency of our approach in context.

5.1 Performance

Figure 6 shows the approximate voxel counts of our test models to give a measure of input complexity. All computations were performed on a personal computer running Windows 10 with an AMD 3900X with 12C/24T(6/64MB L2/3 cache) at stock speeds and 64GB of memory.

All methods were implemented in standard C++17 and multi-threaded using tbb [Intel 2020]. The white noise implementation used was the Mersenne Twister 19937 in the standard C++ library.

Each variation was run 10 times and the median was used in Figure 7.

The standard deviation for each set was below 3% excluding a single outlier run.

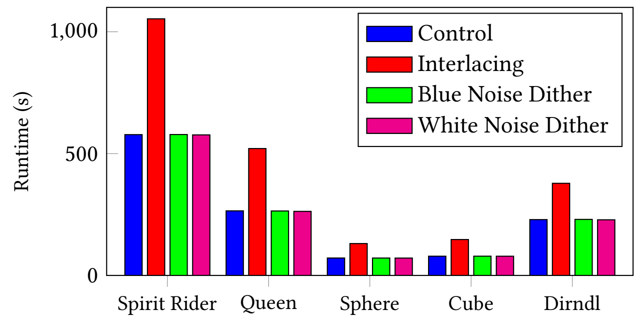

Fig. 7. Runtimes for each model and method.

The voxel counts for the interlacing are double what is given due to the doubling of the Y resolution during computation, which accounts for the nearly double runtime compared to control or dither; the runtime is slightly less than double since file loading and mesh preprocessing time do not depend on the number of voxels. In contrast, dithering incurred almost no extra overhead regardless of the noise used. The runtime for dithering was within the run to run variance of the control. The exact runtimes can be found in the supplementary materials.

5.2 Quantitative Evaluation

In this section, we quantify the advantages of our dithering approach to reduce quantization artifacts via the low-pass filtering effect of the printing process. The true low-pass effect of the printing system is unknown and difficult to measures, although work has been done in this direction [Page et al. 2017]. To approximate this effect, we applied standard low-pass filtering, or smoothing, techniques to meshes generated from the voxel output by marching cubes [Lorensen and Cline. 1987].

We use a cube rotated by 2° about each axis of the printer's build coordinate system as a challenging benchmark object since quantization artifacts are not masked by high frequency surface detail. Figure 8 shows the results, where we evaluate geometric error for all methods. We use the same voxel resolution (d_x, d_y, d_z) that we use for printing in Section 5.3. The point-wise distance to the reference (input) surface provides a measure of the quantization error, which we evaluate before (a) and after (b) applying 160 iterations of Taubin smoothing [Taubin 1995] as implemented in MeshLab [Cignoni et al. 2008]. We chose Taubin smoothing because it makes few assumptions about the geometry and does a good job of preserving volume. We visualize both cumulative error plots and point-wise errors, which are best viewed by zooming in.

Prior to smoothing, dithering with all types of noise introduces additional errors w.r.t. the input surface, as compared to no correction and interlacing, which introduces additional errors primarily along one axis. The cumulative error curves for dithering are virtually identical, regardless of the type of noise used. Note, however, that these errors are decorrelated, whereas those for control and interlacing correlate completely with the printing axes.

After smoothing, the errors for control do not change much and remain aligned with the printing axes. For interlacing, the errors improve primarily along the y-axis, but remain aligned with the

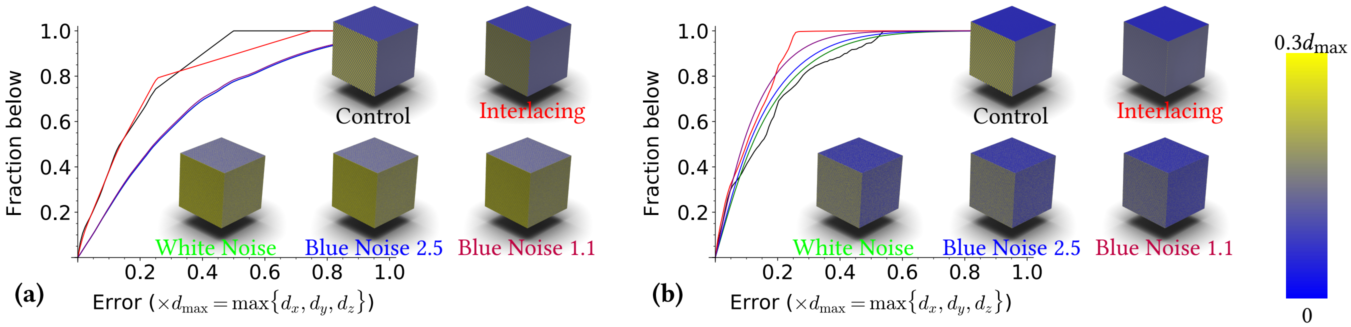

Fig. 8. Simulating the interaction of dithering and interlacing with the low-pass effect of the printing system [Mimaki 2020] used for the results in Section 5.3. We rotated a cube 2° about each of the principle axes, sliced with no staircase correction, interlacing, dithering with white noise, blue noise with $\sigma = 2.5$ and blue noise with $\sigma = 1.1$ (proposed). We then extracted a mesh from the voxel material assignments and measured the error w.r.t. the input mesh (a). To approximate the low-pass effect of the printing system, we applied 160 iterations of Taubin smoothing to the meshes (b), and again measured the error. The plots show cumulative error curves, and the quantization errors are color coded on the surface of the cubes. Best viewed zoomed in. The curve for white noise in (a) overlaps with blue noise 2.5 and can be seen by zooming in.

printing axes. The errors for all types of dithering improve significantly, most so for blue noise with $\sigma = 1.1$, followed by $\sigma = 2.5$ and white noise. These improvements are distributed uniformly over the surface and the remaining errors are much less correlated with the printing axes than for control or interlacing. This validates both the use of 3D blue noise implicit dithering, and tuning of the blue noise mask to $\sigma = 1.1$ per Section 4.2.

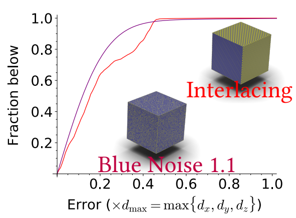

Interlacing is designed for printing systems with a high degree of anisotropy, canceling quantization errors along the lowest-resolution axis. The inset compares interlacing to dithering with an optimized mask on an isotropic grid with $\approx 84\ \mu$m spacing following 160 iterations of smoothing as in Figure 8b, and we see that dithering reduces quantization errors better.

Surface roughness can not only impact the mechanical properties of the printed part [Kampker et al. 2017; Moore and Williams 2015], but also the ease with which support material can be removed. We visually observed the surfaces printed with optimized blue noise dithering to be smoother than the rest, as seen in Section 5.3 and in the supplemental material. We further performed a numerical evaluation of surface roughness using the same simulation process as in this section, which matched these observations. See the supplemental material for details.

We evaluate the perceived quality of the simulated prints assuming the resolution of the Mimaki 3DUJ-553 printer using the SSIM index [Wang et al. 2004]. The ground truth and simulated printer geometry are rendered using Blender's Cycle path tracer under diffuse area illumination from above. The viewing conditions of the renderings are adjusted to a viewing distance of approx. 50cm with a maximum spatial frequency of approx. 50 cycles-per-degree, which is close to the human visual system's sensitivity limit for achromatic contrasts under indoor luminance conditions [Van Nes and Bouman 1967]. Fig.9 shows the SSIM index and maps computed between ground truth geometry and simulated prints. All methods improve the perceived errors of the simulated control print. Blue noise dithering with $\sigma = 1.1$ performs best, followed closely by interlacing, and blue noise dithering with $\sigma = 2.5$. The biggest SSIM errors are created by dithering with white noise.

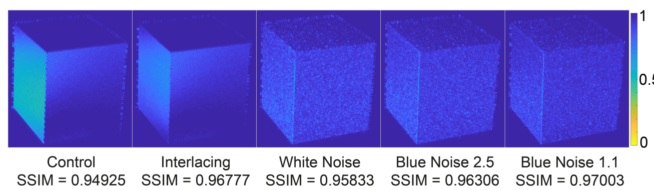

Fig. 9. SSIM maps and indices computed between ground truth and simulated print geometry assuming Mimaki 3DUJ-553 printer resolution.

5.3 Printed results

Figures 1, 10, 11, 12, 13, 14, and 15 show the effect of dithering on printed results. Some areas are highlighted and blown-up, but more artifacts can be observed by zooming in. We observed the pattern that interlacing [Kritchman 2010] removes staircasing artifacts well along the Y-axis, where the resolution is coarsest, but that our dithering approach does a better job of removing quantization errors along the Y-axis, and improves the quality of surfaces of all orientations. Interlacing subsamples the material arrangement produced by 3D halftoning. If a blue-noise halftoning in double resolution is used, such subsampling creates volumetric Moiré resulting in slight color shifts. This can be observed in Fig. 1(b) (hair and skin), Fig. 10(b) (skin) or Fig. 15(b) (base). Halftoning is not effected by implicit shape dithering ensuring a color match with the Control print.

Figures 12 - 14 illustrate the benefits of using an optimized blue-noise mask resulting in apparently smoother surfaces compared to non-optimized blue-noise or white noise. For white noise low-frequency artifacts are particularly disturbing. More comparisons can be found in the supplemental material.

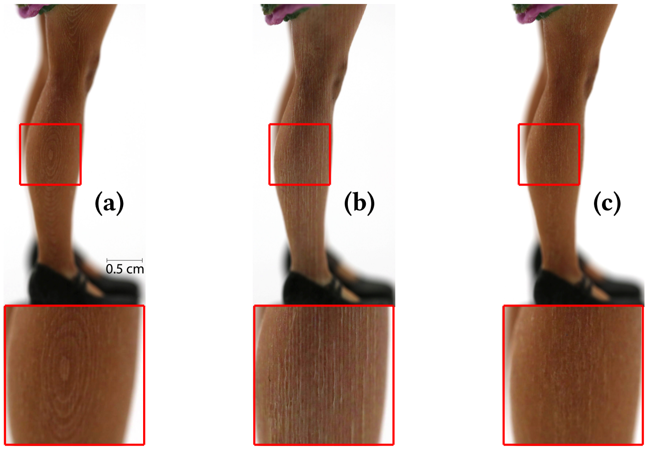

Fig. 10. Quantization artifacts on the legs of the Dirndl model (a) are only partially removed by interlacing (b), but completely removed by dithering (c). More artifacts are visible by zooming in.

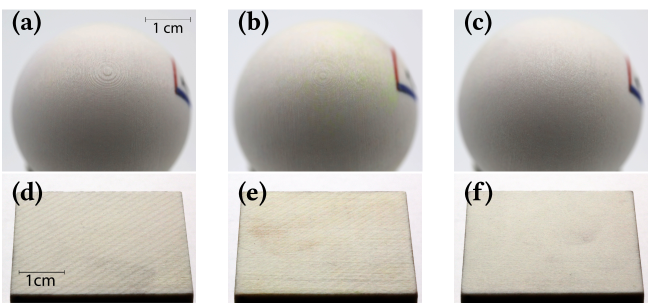

Fig. 11. Quantization artifacts on the sphere (a) are only partially removed by interlacing (b). Dithering (c) removes more artifacts. Similarly for the cube (d-f). Best viewed zoomed in.

6 LIMITATIONS

We optimized the 3D blue noise mask using the reverse generalized Fourier slice theorem, which assumes plane slices through the mask. Object surfaces are curved in general, and the dither signal on curved surfaces may possess larger densities at low-frequencies compared to plane surfaces, particularly if the curvature is large compared to the voxel size. However, in our experiments we could not see any increase of quantization artifacts even on highly curved surfaces. This observation is restricted to the particular printing system and is not generalizable.

In our dithering implementation, we use finite differences to estimate surface normals by computing gradients of the signed distance field. This reduces the resolution of normals, which introduces another type of quantization noticeable by slight low-frequency patterns in renderings of Fig. 8.

Our dithering approach removes quantization-based artifacts, but it cannot inhibit process-based artifacts caused by mechanical leveling mechanisms and mixing of build and support material.

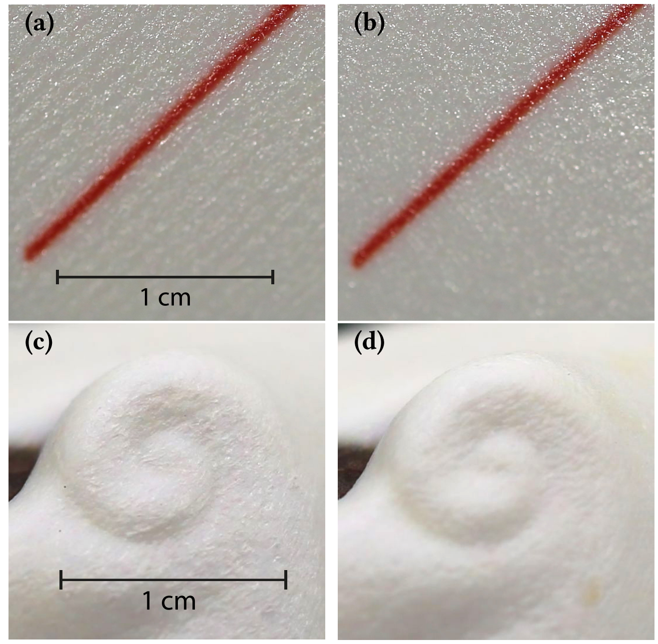

Fig. 12. Low-frequency artifacts remain apparent when using an unoptimized blue noise mask (a) compared to ours (b). The surface smoothness difference can be observed in (c,d). Best viewed zoomed in.

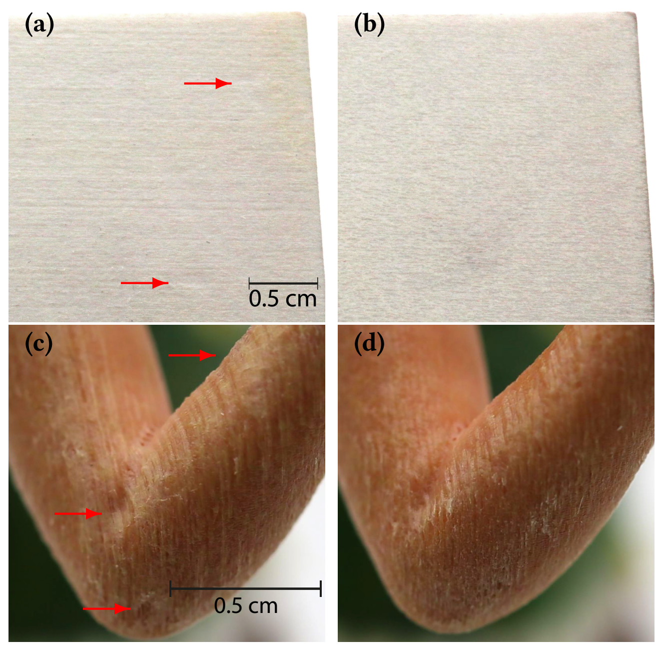

Fig. 13. Using white noise (a,c) for the dither leads to visible artifacts compared to using blue noise (b,d). Best viewed zoomed in.

Fig. 14. The surface smoothness across different noise sources can be observed here. White noise (a), unoptimized blue noise (b), optimized blue noise (c). Best viewed zoomed in.

7 CONCLUSION

We have proposed an efficient dithering approach that improves surface quality and accuracy of 3D prints by removing quantization artifacts caused by discrete voxel sampling of continuous shapes in voxel-controlled 3D printing. Our algorithm is a streaming compatible point process, does not adversely affect color reproduction and comes with negligible computational cost. It displaces the part's implicit surface by a spatially high-frequent signal to shift low-frequent quantization errors to higher spatial frequencies that can be removed by the low-pass filtering mechanisms of the 3D printing process to a large extent. The dither signal is generated by a 3D blue-noise mask optimized to produce minimum low-frequency densities on 2D surfaces by leveraging the reverse generalized Fourier slice theorem. We have verified the geometric accuracy and performance qualitatively and quantitatively. Future work will focus on algorithms to reduce process-based artifacts by extending the dither approach to control build and support material mixing at the part's surface. The goal is to reduce both low-frequency process-based striping artifacts at vertical surfaces and part thickening due to non-removable support material mixed into build material.

ACKNOWLEDGMENTS

This work was funded by the Fraunhofer Internal project Highlight. We would like to thank Dennis Harroun for providing the Spirit Rider model used in the comparisons, as well as Marketiger B.V. for printing all the models.

REFERENCES

3D Printing Industry. 2020. Mimaki 3DUJ-2207 UV-LED 3D Printer. https://3dprintingindustry.com/news/mimaki-opens-up-full-3d-color-printing-for-less-than-e40k-with-3duj-2207-uv-led-3d-printer-178953/.

A.U. Agar and J.P. Allebach. 2005. Model-based color halftoning using direct binary search. *Image Processing, IEEE Transactions on* 14, 12 (2005), 1945–1959.

M. Alexa, K. Hildebrand, and S. Lefebvre. 2017. Optimal Discrete Slicing. *ACM TOG* 36, 1, Article 12 (Jan. 2017), 16 pages. https://doi.org/10.1145/2999536

M. Alexa and J.E. Kyprianidis. 2015. Error diffusion on meshes. *Computers and Graphics (Proc. SMI 2014)* 46 (2015), 336–344.

T. Auzinger, W. Heidrich, and B. Bickel. 2018. Computational design of nanostructural color for additive manufacturing. *ACM TOG (Proc. SIGGRAPH)* 37, 4 (2018), 1–16.

T. Baar, S. Samadzadegan, H. Brettel, P. Urban, and M. V. Ortiz Segovia. 2014. Printing gloss effects in a 2.5 D system. In *IS&T/SPIE Electronic Imaging*. International Society for Optics and Photonics, 90180M–90180M.

V. Babaei, K. Vidimče, M. Foshey, A. Kaspar, P. Didyk, and W. Matusik. 2017. Color contoning for 3D printing. *ACM TOG (Proc. SIGGRAPH)* 36, 124 (2017). Issue 4.

B. E. Bayer. 1973. An optimum method for two-level rendition of continuous-tone pictures. In *IEEE Intl. Conf. on Comm.* Seattle, WA, 11–15.

A. Brunton and L. Abu Rmaileh. 2021. Displaced Signed Distance Fields for Additive Manufacturing. *ACM TOG (Proc. SIGGRAPH)* 40, 4 (2021).

A. Brunton, C. A. Arikan, T. M. Tanksale, and P. Urban. 2018. 3D Printing Spatially Varying Color and Translucency. *ACM TOG (Proc. SIGGRAPH)* 37, 4 (2018), 157:1–157:13.

A. Brunton, C. A. Arikan, and P. Urban. 2015. Pushing the Limits of 3D Color Printing: Error Diffusion with Translucent Materials. *ACM TOG* 35, 1 (2015), 4.

J. Chang, B. Alain, and V. Ostromoukhov. 2009. Structure-Aware Error Diffusion. *ACM TOG (Proc. SIGGRAPH Asia)* 28, 5 (2009), 162:1–162:8.

W. Cho, E.M. Sachs, N. M. Patrikalakis, and D. E. Troxel. 2003. A dithering algorithm for local composition control with three-dimensional printing. *CAD* 35, 9 (2003), 851–867.

Paolo Cignoni, Marco Callieri, Massimiliano Corsini, Matteo Dellepiane, Fabio Ganovelli, and Guido Ranzuglia. 2008. MeshLab: an Open-Source Mesh Processing Tool. In *Eurographics Italian Chapter Conference*, Vittorio Scarano, Rosario De Chiara, and Ugo Erra (Eds.). The Eurographics Association. https://doi.org/10.2312/LocalChapterEvents/ItalChap/ItalianChapConf2008/129-136

D. Cohen-Or and A. Kaufman. 1995. Fundamentals of Surface Voxelization. *Graphical Models and Image Processing* 57, 6 (November 1995), 453–461.

Y. Dong, J. Wang, F. Pellacini, X. Tong, and B. Guo. 2010. Fabricating spatially-varying subsurface scattering. *ACM TOG (Proc. SIGGRAPH)* 29, 4 (2010), 62:1–62:10.

DP Polar. 2020. AMpolar i2. https://www.dppolar.de/en/3d-printer.

O. Elek, D. Sumin, R. Zhang, T. Weyrich, K. Myszkowski, B. Bickel, A. Wilkie, and J. Křivánek. 2017. Scattering-aware Texture Reproduction for 3D Printing. *ACM TOG (Proc. of SIGGRAPH Asia)* 36, 6 (2017), 241:1–241:15.

R.W. Floyd and L. Steinberg. 1976. An adaptive algorithm for spatial grey scale. In *Proceedings of the Society of Information Display*. SID, 75–77.

M. Hašan, M. Fuchs, W. Matusik, H. Pfister, and S. Rusinkiewicz. 2010. Physical reproduction of materials with specified subsurface scattering. *ACM TOG (Proc. SIGGRAPH)* 29, 4 (2010), 61:1–61:9.

Fraunhofer IGD. 2020. Cuttlefish Version 2020.09. https://www.cuttlefish.de/.

Intel. 2020. Intel Threading Building Blocks. https://software.intel.com/content/www/us/en/develop/tools/threading-building-blocks.html.

A. Kampker, K. Kreisköther, and C. Reinders. 2017. Material and Parameter Analysis of the PolyJet Process for Mold Making Using Design of Experiments. *International Journal of Materials and Metallurgical Engineering* 11, 3 (2017), 242 – 249. https://publications.waset.org/vol/123

E. Kritchman. 2010. Method for printing of three-dimensional objects. US Patent 7,658,976.

Ares Lagae and George Drettakis. 2011. Filtering solid Gabor noise. *ACM Transactions on Graphics (TOG)* 30, 4 (2011), 1–6.

Y. Lan, Y. Dong, F. Pellacini, and X. Tong. 2013. Bi-scale appearance fabrication. *ACM TOG (Proc. SIGGRAPH)* 32, 4 (2013), 145–1.

D.L. Lau, G.R. Arce, and N.C. Gallagher. 1999. Digital halftoning by means of green-noise masks. *JOSA A* 16, 7 (1999), 1575–1586.

D. L. Lau and G. R. Arce. 2001. *Modern digital halftoning*. CRC Press.

W.E. Lorensen and H.E. Cline. 1987. Marching Cubes: A High Resolution 3D Surface Construction Algorithm. *SIGGRAPH Comput. Graph.* 21, 4 (Aug. 1987), 163–169. https://doi.org/10.1145/37402.37422

A. Luongo, V. Falster, M.B. Doest, M.M. Ribo, E.R. Eiríksson, D.B. Pedersen, and J.R. Frisvad. 2020. Microstructure control in 3D printing with digital light processing. 39, 1 (2020), 347–359.

T. Malzbender, R. Samadani, S. Scher, A. Crume, D. Dunn, and J. Davis. 2012. Printing reflectance functions. *ACM TOG* 31, 3 (2012), 1–11.

W. Matusik, B. Ajdin, J. Gu, J. Lawrence, H. Lensch, F. Pellacini, and S. Rusinkiewicz. 2009. Printing spatially-varying reflectance. In *ACM Transactions on Graphics (TOG)*, Vol. 28. ACM, 128.

Mimaki. 2020. 3DUJ-553 3D Printer. https://www.mimakieurope.com/products/3d/3duj-553/.

T. Mitsa and K.J. Parker. 1992. Digital halftoning technique using a blue-noise mask. *JOSA A* 9, 11 (1992), 1920–1929.

J.P. Moore and C.B. Williams. 2015. Fatigue properties of parts printed by PolyJet material jetting. *Rapid Prototyping Journal* (2015).

P. Morovič, J. Morovič, J. Gondek, and R. Ulichney. 2017a. Direct pattern control halftoning of Neugebauer primaries. *IEEE TIP* 26, 9 (2017), 4404–4413.

P. Morovič, J. Morovič, I. Tastl, M. Gottwals, and G. Dispoto. 2017b. HANS3D: a multi-material, volumetric, voxel-by-voxel content processing pipeline for color and beyond. In *Color and Imaging Conference*, Vol. 2017. Society for Imaging Science and Technology, 219–225.

E. Napadensky. 2014. Method and system for three-dimensional fabrication. US Patent 8,784,723.

R. Ng. 2005. Fourier slice photography. *ACM TOG* (2005), 735–744.

A. Orth, K.L. Sampson, K. Ting, J. Boisvert, and C. Paquet. 2021. Correcting ray distortion in tomographic additive manufacturing. *Optics Express* 29, 7 (Mar 2021), 11037–11054. https://doi.org/10.1364/OE.419795

V. Ostromoukhov. 2001. A Simple and Efficient Error-Diffusion Algorithm. In *Proc. SIGGRAPH*.

M. Page, G. Obein, C. Boust, and A. Razet. 2017. Adapted modulation transfer function method for characterization and improvement of 3D printed surfaces. *Electronic Imaging* 2017, 8 (2017), 92–100.

W.-M. Pang, Y. Qu, T.-T. Wong, D. Cohen-Or, and P.A. Heng. 2008. Structure-aware halftoning. *ACM TOG (Proc. SIGGRAPH)* 27, 3 (2008), 89.

C. Peters. 2017. Moments in Graphics: The problem with 3D blue noise. http://momentsingraphics.de/3DBlueNoise.html.

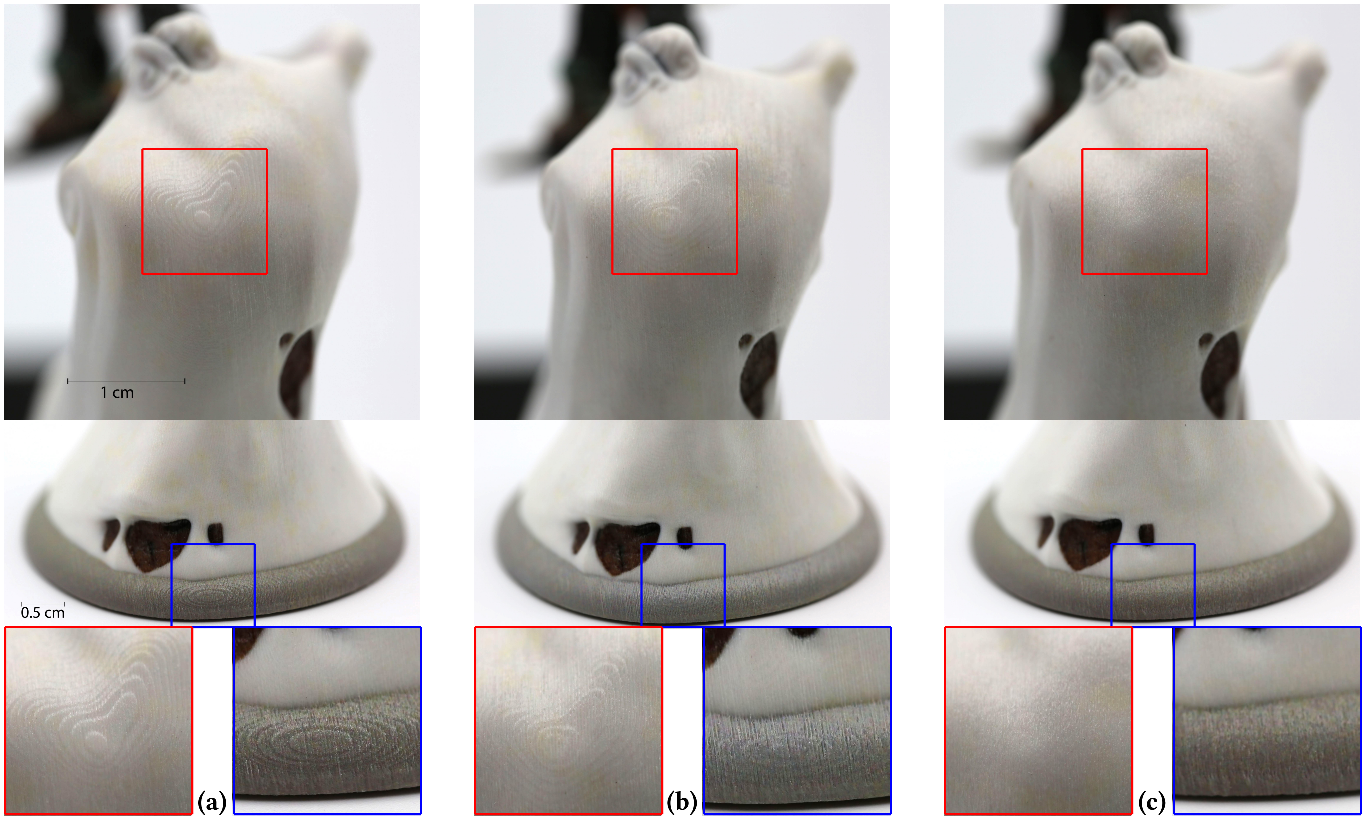

Fig. 15. Quantization artifacts on smooth surfaces of the snow and base of the Spirit Rider model (a), are only partially removed by interlacing (b). Our blue noise implicit dither approach removes all of them. More artifacts are visible by zooming in.

M. Piovarči, M. Foshey, V. Babaei, S. Rusinkiewicz, W. Matusik, and P. Didyk. 2020. Towards spatially varying gloss reproduction for 3D printing. *ACM TOG (Proc. SIGGRAPH Asia)* 39, 6 (2020), 1–13.

Quantica. 2020. Tech Platform. https://quantica3d.com/technology/.

O. Rouiller, B. Bickel, J. Kautz, W. Matusik, and M. Alexa. 2013. 3D-printing spatially varying BRDFs. *IEEE CG&A* 33, 6 (2013), 48–57.

S. Samadzadegan, T. Baar, P. Urban, M.V.O. Segovia, and J. Blahová. 2015. Controlling colour-printed gloss by varnish-halftones. In *Measuring, Modeling, and Reproducing Material Appearance 2015*, Vol. 9398. International Society for Optics and Photonics, 93980V.

Stratasys. 2020. J8 Series 3D Printers. https://www.stratasys.com/3d-printers/j8-series.

D. Sumin, T. Rittig, Babaei V, T. Nindel, A. Wilkie, P. Didyk, B. Bickel, J. KŘivánek, K. Myszkowski, and T. Weyrich. 2019. Geometry-aware scattering compensation for 3D printing. *ACM TOG (Proc. SIGGRAPH)* 38, 4 (2019).

G. Taubin. 1995. A Signal Processing Approach to Fair Surface Design. In *Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '95)*. Association for Computing Machinery, New York, NY, USA, 351–358. https://doi.org/10.1145/218380.218473

R.A. Ulichney. 1993. Void-and-cluster method for dither array generation. In *Human Vision, Visual Processing, and Digital Display IV*, Vol. 1913. International Society for Optics and Photonics, 332–343.

F. L. Van Nes and M. A. Bouman. 1967. Spatial modulation transfer in the human eye. *JOSA* 57, 3 (1967), 401–406.

Z. Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli. 2004. Image Quality Assessment: From Error Visibility to Structural Similarity. *IEEE Transactions on Image Processing* 13, 4 (2004), 600–612.

Xaar. 2020. High Laydown Technology. https://www.xaar.com/en/about/xaar-technologies/high-laydown-technology/.

D.-M. Yan, J.-W. Guo, B. Wang, X.-P. Zhang, and P. Wonka. 2015. A Survey of Blue-Noise Sampling and Its Applications. *J. Comp. Sci. and Tech.* 30 (2015), 439–453.

B. Zhou and X. Fang. 2003. Improving Mid-tone Quality of Variable-Coefficient Error Diffusion Using Threshold Modulation. *ACM TOG (Proc. SIGGRAPH)* 22, 3 (2003), 437–444.

A PROOF OF COROLLARY 1

We use the same notation and definitions as in [Ng 2005]:

Integral Projection: The canonical projection operator that reduces an N-dimensional function f down to M dimensions by integrating over the first dimension is denoted as I_M^N and defined as $I_M^N[f](x_1,\dots,x_M) = \int f(x_1,\dots,x_N)\ dx_{M+1}\dots dx_N$.

Slicing Transform: The canonical slicing operator that reduces an N-dimensional function f down to M dimensions is denoted as S_M^N and defined as $S_M^N[f](x_1,\dots,x_M) = f(x_1,\dots,x_M,0,\dots,0)$.

Fourier Transform: The N dimensional Fourier transform operator is denoted by F^N and its inverse by F^{-N}.

Basis Change: Basis change is performed by an NxN-dimensional invertible matrix B using the following notation $B[f](x_1,\dots,x_N) = f(B^{-1}(x_1,\dots,x_N)^T)$.

Corollary 1

$$F^M \circ S_M^N \circ B = I_M^N \circ \left(B^{-T}/|B^{-T}|\right) \circ F^N \tag{12}$$

Proof: Ng [2005] showed that basis change and Fourier transform commute as follows:

$$F^N \circ B = \left(B^{-T}/|B^{-T}|\right) \circ F^N. \tag{13}$$

Thus, it is sufficient to show that

$$I_M^N \circ F^N = F^M \circ S_M^N \tag{14}$$

because then

$$I_M^N \circ \left(B^{-T}/|B^{-T}|\right) \circ F^N \overset{(13)}{=} I_M^N \circ F^N \circ B \overset{(14)}{=} F^M \circ S_M^N \circ B. \quad (15)$$

To show (14), let $G(u_1,\ldots,u_N) = F^N[g](u_1,\ldots,u_N)$ be the Fourier transform of g, then

$$\begin{aligned}(I_M^N \circ F^N)[g](u_1,\ldots,u_N) &= I_M^N \circ G(u_1,\ldots,u_N)\\ &= \int G(u_1,\ldots,u_N)\, du_{M+1}\ldots du_N \end{aligned} \quad (16)$$

and

$$\begin{aligned}
&S_M^N[g](x_1,\ldots,x_M)\\
&= (S_M^N \circ F^{-N})[G](x_1,\ldots,x_M)\\
&= S_M^N \circ \int G(u_1,\ldots,u_N)\exp(2\pi i \sum_{i=1}^N u_i x_i) du_1\ldots du_N\\
&= \int G(u_1,\ldots,u_N)\exp(2\pi i \sum_{i=1}^M u_i x_i) du_1\ldots du_N\\
&= \int \left(\int G(u_1,\ldots,u_N) du_{M+1}\ldots du_N\right)\exp(2\pi i \sum_{i=1}^M u_i x_i) du_1\ldots du_M\\
&\overset{(16)}{=} \int \left(\left(I_M^N \circ F^N[g]\right)(u_1,\ldots,u_N)\right)\exp(2\pi i \sum_{i=1}^M u_i x_i) du_1\ldots du_M\\
&= \left(F^{-M} \circ I_M^N \circ F^N\right)[g](x_1,\ldots,x_M).
\end{aligned} \quad (17)$$

Thus

$$I_M^N \circ F^N = F^M \circ F^{-M} \circ I_M^N \circ F^N \overset{(17)}{=} F^M \circ S_M^N. \quad (18)$$

□