Stochastic Cloning and Smoothing for Fusion of Multiple Relative and Absolute Measurements for Localization and Mapping

Thomas Emter and Janko Petereit
Fraunhofer Institute for Optronics, System Technologies and Image Exploitation IOSB,
Karlsruhe, Germany.
Email: thomas.emter@iosb.fraunhofer.de and janko.petereit@iosb.fraunhofer.de

***Abstract*— A mobile robot is reliant on precise and robust localization and mapping for autonomous navigation. For this purpose, sensor fusion techniques are employed to combine measurements of multiple sensor data sources. The well-known Extended Kalman filter is the standard approach to integrate absolute measurements; however, multiple relative measurements, i.e., measured differences between the current system state and a past system state, cannot be directly incorporated into the filter. This paper presents a fusion algorithm for the integration of absolute and multiple relative measurements for localization and mapping of mobile robots. A novel approach exploiting concurrent stochastic cloning and smoothing is introduced for robust inclusion of additional relative measurements. The proposed fusion method is applied to perform simultaneous localization and mapping with sensor data from an IMU, a GPS, wheel odometry, and scan matching of data from a 3D LiDAR.**

***Index Terms*—Navigation, localization, mapping, SLAM, stochastic cloning, smoothing, loop closure.**

I. INTRODUCTION

Autonomous navigation for mobile robots is reliant on precise self-localization based on sensor measurements. Since these measurements are subject to noise and errors, multiple sensors may be combined by means of sensor fusion in order to improve both precision and robustness.

An Extended Kalman filter (EKF) is a well-known and effective fusion algorithm. In its basic form, the EKF is capable of fusing a system model or at most one relative sensor with multiple absolute sensors; measurements of additional relative sensors cannot be integrated directly.

Precision and robustness of the localization can be further improved with perception sensors for simultaneous mapping. Therefore, algorithms for simultaneous localization and mapping (SLAM) incorporate methods to model the dependencies between the localization and the map. SLAM algorithms may also include loop closure handling to enhance the estimation.

In this paper, a novel fusion algorithm based on the EKF is presented. An inertial measurement unit (IMU) sensor, a GPS, and a wheel odometry sensor are combined with scan matching of 3D LiDAR data for simultaneous localization and mapping. Therefore, additional relative measurements are integrated into the fusion algorithm by employing an augmentation scheme. A supplementary smoothing step is introduced to the augmentation for a stable filtering.

The following section covers related work. Afterwards, the proposed sensor fusion method is described and the mapping scheme is explained. Thereafter, results are discussed and finally a conclusion closes the paper.

II. RELATED WORK

For localization, Extended Kalman filters have been widely used for sensor fusion, e.g. [1] and [2]. In most systems the sensors are not synchronized, which means their measurements have to be processed asynchronously. Therefore, Welch et al. proposed a fusion framework based on an EKF capable of processing measurements from asynchronous sensors [3]. If multiple relative sensors are present, the standard EKF formulation is not sufficient anymore because past states become correlated with the current state and an augmented state estimation is necessary. Roumeliotis et al. introduced *Stochastic Cloning* for augmentation by cloning the respective states connected by a relative measurement [4]. The stochastic cloning method was included in an EKF-based framework with explicit handling of asynchronous sensors by interpolating the respective proprioceptive measurements to get the best available linearization point by Lynen et al. [5].

In [6] it is shown that if the variance of the heading remains small, EKF-based 2D SLAM is robust concerning loops closures and consistent maps can be built. For the *FastSLAM 2.0* algorithm, convergence was proven with only one particle which is similar to EKF-SLAM [7]. The paper presenting the *Hector SLAM* algorithm additionally shows that no explicit loop closure may be necessary if localization is very precise [8]. When additionally incorporating absolute sensors like GPS, loop closure is even more robust, as shown in [9], where a FastSLAM 2.0 algorithm combined with a grid map successfully closes a large loop with only one particle.

3D SLAM is much more elaborate due to the increased degrees of freedom and memory requirements. For 3D mapping, often graph-based methods are used [10]; however, achieving real-time performance is difficult and requires a combination with filtering [11].

In order to estimate the relative pose between scans from a 3D LiDAR sensor, scan matching algorithms have been developed. In the *Velodyne SLAM* algorithm, surfaces are extracted from the scans to efficiently estimate the poses, which makes it well suited for urban scenarios [12]. Point based methods like *iterative closest point* (ICP) make no assumptions on the scenario and thus are better suited

for off-road or heterogeneous environments [13]. A variant especially developed for scanning sensors is the *Generalized-ICP* (GICP) algorithm [14].

III. SENSOR FUSION

For localization, the full 6DoF pose is estimated. The fusion framework is based on an EKF and all measurements can be processed asynchronously. In order to avoid local discontinuities due to geodetic projections, the global position is estimated in latitude, longitude, and altitude. Attitude and heading is represented by a quaternion. The filtering state is formulated in the error state space:

$$\tilde{\mathbf{x}}_{\text{base}} = [\Delta\mathbf{p}\ \Delta\boldsymbol{\Theta}\ \Delta\mathbf{v}\ \Delta\mathbf{b}_{\omega}\ \Delta\mathbf{b}_{\text{acc}}\ \Delta s_{\text{odo}}]^{\mathrm{T}} \tag{1}$$

with the metric errors of the 3D position, 3D angular errors, 3D velocity errors, bias errors of the gyroscopes and accelerometers of the IMU, and a scale factor error of the wheel odometer. The error state space estimation has the advantage that although the position is estimated in latitude and longitude, the correction can be processed in a local metric space. The same applies to the attitude and heading quaternion, which can be corrected with angular deltas.

The prediction is computed with a strapdown algorithm from high-rate IMU measurements. The accelerometers of the IMU are used for attitude stabilization, and measurements of the GPS are integrated as absolute position and velocity updates. The yaw angle is supported by a heading estimator from past GPS position measurements and also updated by a non-holonomic constraint based on the kinematic model of the mobile robot. The wheel odometry measures the forward velocity of the vehicle. For scan matching, the GICP algorithm [14] is used on inertially corrected scans. The integration of the relative updates of the scan matching can be interpreted as LiDAR odometry and is described in the following paragraph.

A. Stochastic Cloning EKF

The integration of relative measurements from scan matching into the EKF is not straightforward. The EKF is based on the Markov assumption, which means the current state incorporates all information about the processed measurements and past states. A relative measurement introduces dependencies between states which violates the Markov assumption, and for its integration an augmentation of the state vector is necessary.

For augmentation, a stochastic cloning update is employed according to [4] and [5]. Assuming a relative measurement $\mathbf{z}_{k,k+m}$ between the two states at time k and $k+m$ respectively the augmented state and covariance matrix become:

$$\breve{\mathbf{x}}_{k+m|k} = [\hat{\mathbf{x}}_{k|k}^{\mathrm{T}}\ \ \hat{\mathbf{x}}_{k+m|k}^{\mathrm{T}}]^{\mathrm{T}} \text{ and} \tag{2}$$

$$\breve{\mathbf{P}}_{k+m|k} = \begin{bmatrix} \mathbf{P}_{k|k} & \mathbf{P}_{k|k}\boldsymbol{\mathcal{F}}_{k+m,k}^{\mathrm{T}} \\ \boldsymbol{\mathcal{F}}_{k+m,k}\mathbf{P}_{k|k} & \mathbf{P}_{k+m|k} \end{bmatrix} \tag{3}$$

with the cumulative product of the system dynamic matrixes $\boldsymbol{\mathcal{F}}_{k+m,k} = \prod_{i=1}^{m} \mathbf{F}_{k+i}$.

The measurement residual is

$$\mathbf{y}_{k+m} = \mathbf{z}_{k,k+m} - \hat{\mathbf{z}}_{k,k+m} \tag{4}$$

where $\hat{\mathbf{z}}_{k,k+m}$ is the estimated difference between state $\mathbf{x}_k$ and $\mathbf{x}_{k+m}$.

The Kalman gain is computed by

$$\breve{\mathbf{K}} = \breve{\mathbf{P}}_{k+m|k}\breve{\mathbf{H}}^{\mathrm{T}}\breve{\mathbf{S}}_{k+m}^{-1} \tag{5}$$

with

$$\breve{\mathbf{S}}_{k+m} = \breve{\mathbf{H}}\breve{\mathbf{P}}_{k+m|k}\breve{\mathbf{H}}^{\mathrm{T}} + \mathbf{R}_{k,k+m} \tag{6}$$

where $\mathbf{R}_{k,k+m}$ is the covariance of the scan matching. $\breve{\mathbf{H}} = [\mathbf{H}_{k|k}\ \ \mathbf{H}_{k+m|k}]$ contains both measurement Jacobians related to the respective states $\mathbf{x}_k$ and $\mathbf{x}_{k+m}$.

Finally the new estimate and covariance matrix are given by

$$\hat{\mathbf{x}}_{k+m|k+m} = \hat{\mathbf{x}}_{k+m|k} + \mathbf{K}_{k+m}\mathbf{y}_{k+m} \quad \text{and} \tag{7}$$

$$\mathbf{P}_{k+m|k+m} = \mathbf{P}_{k+m|k} - \mathbf{K}_{k+m}\breve{\mathbf{S}}_{k+m}\mathbf{K}_{k+m}^{\mathrm{T}}, \tag{8}$$

where $\mathbf{K}_{k+m}$ can be retrieved from

$$\breve{\mathbf{K}} = [\mathbf{K}_k^{\mathrm{T}}\ \ \mathbf{K}_{k+m}^{\mathrm{T}}]^{\mathrm{T}}. \tag{9}$$

B. Smoothing

When absolute measurements from GPS are available, a severe problem may arise if they are not synchronized to the relative updates of scan matching. If a GPS measurement at time t_{k+n} is incorporated as a filter update, it will introduce an absolute correction. However, a following relative measurement update at time t_{k+m} with $m > n$ may be related to a filter state t_k, which did not get the correction by the absolute measurement at time t_{k+n}. This means the estimates corrected only by the relative measurements are subject to drift despite the presence of an absolute sensor. If no GPS measurement updates the filter at the exact same time as a relative measurement, the error of theses updates is unbounded. Because the relative measurements obtained from scan matching are very accurate, this will lead to a jagged and diverging path and ultimately to a degenerating filter, cf. Fig. 1.

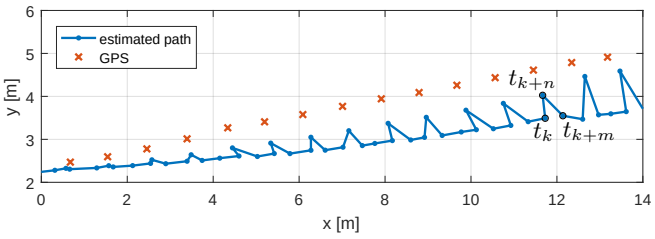

Fig. 1. Jagged path due to asynchronous absolute measurements.

Thus, the absolute updates of GPS should also impact previous updates from relative measurements. Therefore, a backpropagation of the current absolute localization information to earlier estimates by applying a smoothing scheme to the filtered estimates is proposed. For a new relative

measurement at time t_{k+m}, the estimates are smoothed in reverse order up until its related state at time t_k.

The smoothing is performed with a Rauch-Tung-Striebel smoother [15]. For the smoothing scheme, the predicted state and covariance matrix of the standard Kalman filter

$$\hat{\mathbf{x}}_{k+1|k} = \mathbf{F}_{k+1}\hat{\mathbf{x}}_{k|k} \quad \text{and} \tag{10}$$

$$\mathbf{P}_{k+1|k} = \mathbf{F}_{k+1}\mathbf{P}_{k|k}\mathbf{F}_{k+1}^{\mathrm{T}} + \mathbf{Q}_k \tag{11}$$

are fed into the backward recursion yielding the smoothed state and covariance matrix denoted with a superscripted s:

$$\mathbf{A}_k = \mathbf{P}_k\mathbf{F}_{k+1}^{\mathrm{T}}\mathbf{P}_{k+1|k}^{-1}, \tag{12}$$

$$\hat{\mathbf{x}}_k^{\mathrm{s}} = \hat{\mathbf{x}}_k + \mathbf{A}_k(\hat{\mathbf{x}}_{k+1}^{\mathrm{s}} - \hat{\mathbf{x}}_{k+1|k}) \quad \text{and} \tag{13}$$

$$\mathbf{P}_k^{\mathrm{s}} = \mathbf{P}_k + \mathbf{A}_k(\mathbf{P}_{k+1}^{\mathrm{s}} - \mathbf{P}_{k+1|k})\mathbf{A}_k^{\mathrm{T}}. \tag{14}$$

Upon receiving a relative measurement between states at time t_k and t_{k+m}, smoothing is performed feeding back information to state $\hat{\mathbf{x}}_k$. Afterwards, the stochastic cloning update is based on $\hat{\mathbf{x}}_k^{\mathrm{s}}$, thus also utilizing newer information up to time t_{k+m}. In contrast to a fixed-lag smoother, this scheme does not introduce any delay, because the smoothing is not reliant on future measurements or estimates.

Fig. 2 shows how the smoothing feeds back the information to earlier states and prevents the path from diverging as seen before in Fig. 1. Please note, that in Fig. 1 to Fig. 2 for clarity reasons, the estimates of the updates by GPS and scan matching are plotted only as the predictions and updates by the IMU have a tenfold rate.

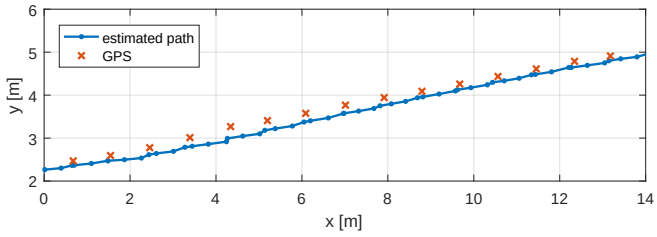

Fig. 2. Path with smoothing.

In Lynen et al. [5] filter divergence does not occur as their realization adds a 6DoF offset to the state accounting for the drift between the world frame and the frame the relative sensor measurements are expressed in. With our proposed smoothing method, the estimation of an additional 6DoF offset is not necessary and the relative motion can be included directly.

IV. MAPPING

A map is not built explicitly on-line but by subsequently performing an implicit mapping process incorporating the localization with relative scan matching updates as explained in Section III. A map can be built after a completed exploration tour or continuously with intermittent map calculations and updates.

A. Map Format

For mapping a Normal Distributions Transform (NDT) voxel map is chosen [16]. An NDT map is built by tessellating the environment into a voxel grid. For the points inside each voxel, the respective normal distributions are estimated. This allows for a compact representation while still preserving shape information.

To correct errors and cope with dynamic objects, which should not be integrated into the map, a method for taking negative information into account is included. An existence probability is attached to each voxel, which is increased by each observation. While tracing the laser ray through free space in the voxel map with a 3D Bresenham algorithm [17], existence probability per voxel is decreased. Voxels with an existence probability falling below a certain threshold are eliminated.

B. Final Path Smoothing

By applying the aforementioned smoothing for the entire path a further refinement of the estimated path can be achieved for mapping, cf. Fig. 3 compared to Fig. 2. Of course, the final smoothing has no impact on the precision of the on-line localization but enhances the mapping by reducing jumps in the localization.

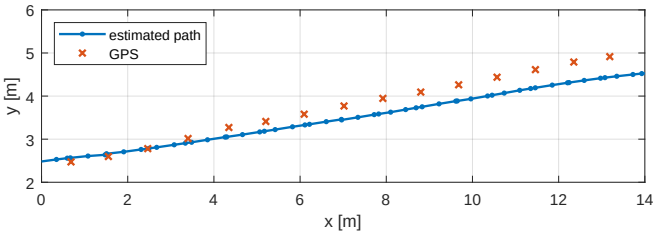

Fig. 3. Path with final smoothing over the entire path.

C. Loop Closure Handling

In modern GPS receivers, most error sources can be mitigated, especially with correction services like SBAS. The remaining errors are dominated by multipath phenomenons, which may lead to deviations up to several meters [18]. The multipath situation is dependent both on the environment and on the satellite constellation. Due to the latter varying over time the errors may change between visits of a position, i.e. loop closures.

The GICP matching between consecutive scans can also be used for loop closure. For each current scan, a check for a potential loop closure is performed backwards in time by calculating the horizontal Euclidean distance of the respective poses. Afterwards, the GICP algorithm is executed and the fitness score of the matching is additionally used to verify or reject the occurrence of a loop closure. The relative motion estimated by the GICP cannot be used as a direct update via stochastic cloning because due to the long time between the two states the assumption of linearization of $\mathcal{F}$ does not hold true and would lead to instabilities of the filter.

The GPS, being the only absolute sensor concerning the position, has a high impact on the estimated position. Therefore, an additional GPS offset is introduced to the continuously estimated state vector to allow for more robust loop closure. It is sufficient to include the position offsets because the attitude and heading are reasonably well stabilized by accelerometer and GPS heading updates. The error state space of Equation 1 expanded with the 3D GPS offset error becomes

$$\tilde{\mathbf{x}}_{\text{loop}} = [\tilde{\mathbf{x}}_{\text{base}}^{\text{T}} \; \Delta\mathbf{o}_{\text{GPS}}]^{\text{T}}. \tag{15}$$

V. RESULTS

In the following results, qualitative evaluations and quantitative comparisons are presented because no ground truth is available. The IMU and GPS measurements are from an *Xsens MTi-G-700* and the LiDAR data is from a *Velodyne HDL-64E*. The *Xsens* has rates of 100 Hz and 4 Hz for its IMU and GPS respectively. The wheel odometry sensor is custom-made and has a rate of 10 Hz. The Velodyne has a scanning rate of 10 Hz, i.e., a rotation takes 100 ms. The vehicle moves during this time leading to a skewing of the scan; hence, the resulting scan has to be rectified. This is accomplished by inertial correction based on the high-rate 6DoF pose estimation of the EKF.

A. Impact of Scan Matching

Using only the raw measurements of the IMU, GPS, and wheel odometry, the heading is not observable. As mentioned before, the heading is supported by an estimator from past GPS positions. Although this provides absolute stabilization to the heading, there are still inaccuracies in certain areas, especially in areas where the vehicle is exposed to quick turns. Fig. 4(a) shows a part of the map where inaccuracies in the heading angle lead to skewed and multiply mapped walls.

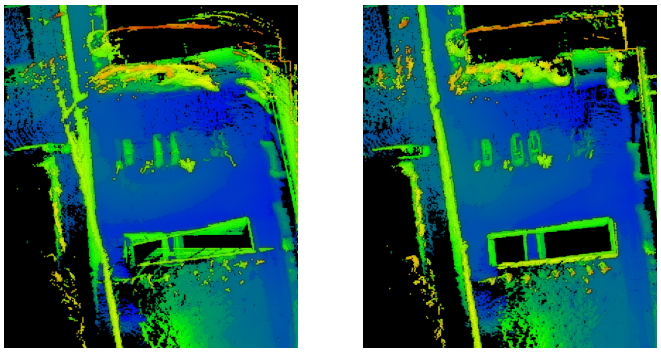

(a) (b)

Fig. 4. Mapping with IMU, GPS, and wheel odometry only (a) compared to including additional scan matching (b). The height is color coded from blue (low) via green to red (high).

Additional scan matching integrated by stochastic cloning increases the accuracy of the estimates and especially the heading is stabilized much more robustly. The resulting map is very precise, only few areas with small visible inaccuracies remain. Fig. 4(b) shows the same area as Fig. 4(a) and the higher precision can be clearly seen.

B. Simulated GPS Outage

In order to evaluate the robustness of the fusion algorithm against failure of an absolute sensor, an outage of the GPS for about 35 seconds or 130.7 m is simulated. Fig. 5 shows the fusion of IMU, GPS, and wheel odometry (purple) compared to the fusion with additionally included scan matching (green). The gray area depicts the interval of the simulated GPS outage. The position error is evaluated with reference to the respective fusion without the outage. Hence, the error is zero before the outage occurs and converges to zero after the outage. The maximum error is nearly four times higher without scan matching but the errors of both paths converge to zero, showing the robustness of the fusion framework against sensor outages.

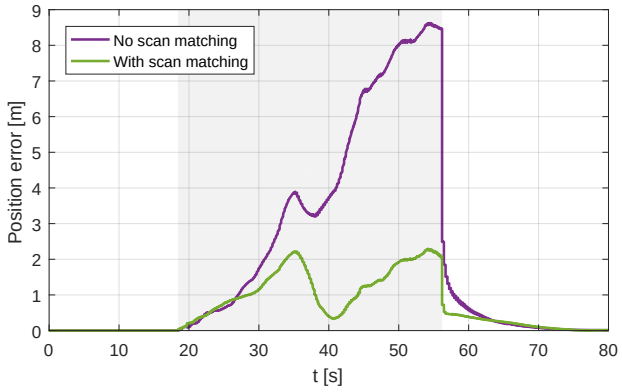

Fig. 5. Simulated GPS Outage (during gray area).

Fig. 6 shows the same simulated GPS outage, this time, the entire path is subsequently smoothed for mapping as explained in Section IV-B. The maximum error of the path without scan matching (purple) is reduced to less than half compared to Fig. 5. The maximum error of path with scan matching (green) is not reduced but the steep correction at the time the GPS outage ended is avoided, which enhances the final mapping.

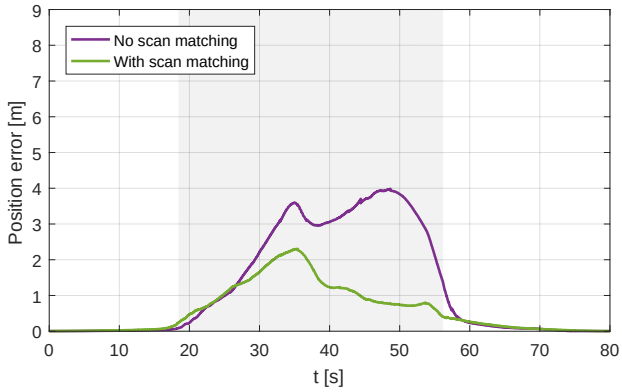

Fig. 6. Simulated GPS Outage (during gray area) with final smoothing.

C. Loop Closure

The altitude above the earths' ellipsoid measured by GPS is much more inaccurate compared to latitude and

longitude [19]. Thus, the performance of the loop closure is demonstrated with the estimated altitude in the following.

Fig. 7 shows a section of the path where a loop closure occurred in a flat area, which means the altitude should be nearly constant. The circle marks the beginning of the tour. In this evaluation, the vehicle is traveling to the north and coming back from the south after a circle of about 550 meters and continuing on the same path for a second turn. The altitude is plotted over the UTM northing; the UTM easting of the path is nearly constant in this section. Without the loop

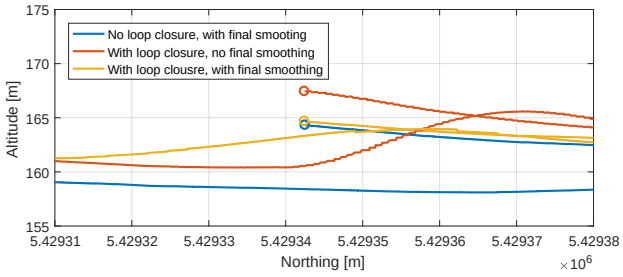

Fig. 7. Altitude in the region of a loop closure.

closure, a difference of about 5 meters can be observed (blue path). Employing the loop closure via estimation of a GPS offset the red and yellow paths clearly show how the path is pulled towards the previous path upon loop closure. While the red path, where no final smoothing is performed, needs a long time to reach the previous path and still does not match it completely. Performing the final smoothing improves both the time to approach the previous path and the alignment. Although errors along the path may be several meters due to poor GPS quality, the error can be decreased to well below one meter in the areas of a loop closure. The remaining error in the altitude may be further reduced especially along the path by including a pressure sensor [20].

D. Complete Map

Fig. 8 shows the final NDT map over a satellite image and the superimposed path in blue. The altitude is shown from blue (low) via green to red (high). A path of about 1.8 km has been covered covering both urban and unstructured off-road scenery with multiple large and small loop closures. The final smoothing further enhances the precision for mapping, because it is possible to also correct backwards, which would not be possible with pure filtering.

E. Timing

To achieve real-time performance and be able to use the filter for on-line localization of a mobile robot, the processing of the measurements has to fulfill specific timing requirements. First, the time to integrate each measurements has to be below the respective sensor rate and below a maximum admissible delay required by algorithms utilizing the localization like path planning and control or safety features. The following results have been measured on an i7-7600U CPU, a mobile processor with two cores running at 2.8 GHz, and 16 GB RAM.

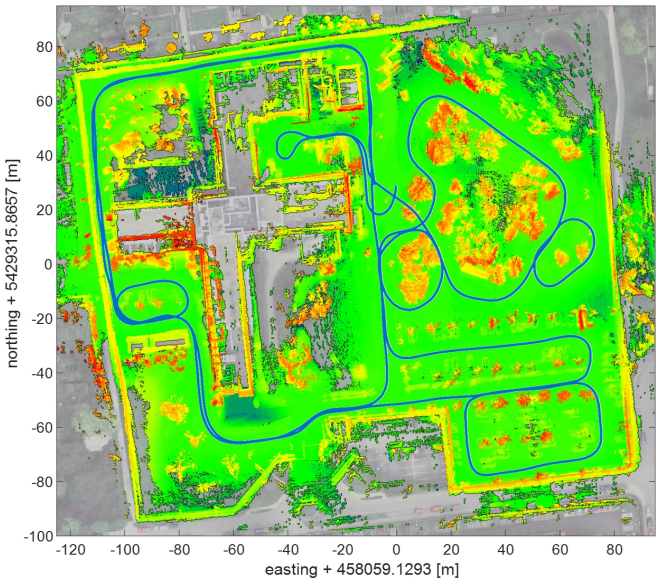

Fig. 8. NDT map over satellite image and superimposed path in blue. The height is color coded from blue (low) via green to red (high). UTM coordinates in zone 32U.

TABLE I
TIMINGS OF EKF UPDATES IN µs

		$t_{\min}$	t_{avg}	$t_{\max}$	σ_t
IMU	@ 100 Hz	63.0	84.9	303.0	18.9
GPS	@ 4 Hz	13.0	55.5	152.0	12.9
Odometry	@ 10 Hz	19.0	31.1	176.0	7.0
LiDAR	@ 10 Hz	259.0	336.1	442.0	23.5

The amounts of time to integrate the measurements into the estimation are listed in Table I. The IMU measurements are used to process the strapdown algorithm for prediction, an attitude update. The IMU processing also includes the update with a non-holonomic constraint. The processing of the GPS measurements includes the position and velocity updates plus the heading estimator. Wheel odometry only involves a normal filter update. The integration of retaliative motion derived from the 3D LiDAR comprises the smoothing with the stochastic cloning update. Even the $t_{\max}$ of each data type and processing is well below their inverse measuring rate and combining all sensor data of one second would take below 35 ms in the worst case. The maximum delay for a single update is below 0.5 ms, which is more than sufficient for path planning and control. The timing for GICP

TABLE II
TIMINGS OF GICP MATCHING IN ms

	$t_{\min}$	t_{avg}	$t_{\max}$	σ_t
GICP	84.1	265.8	567.5	106.4

scan matching is listed in Table II. The average processing time of more than 250 ms is above the 100 ms interval of measurement updates. Therefore, scan matching cannot be processed for every single pair of scans in real-time.

Although omitting pairs of scans for matching in principle allows for real-time processing, a severe delay of up to more than half a second would be too much for autonomous path planning and safe navigation. Therefore, a parallel filter structure is implemented, which is explained in detail in [21]. Two filters are running in parallel. While an on-line EKF processes only GPS and IMU data, a mapping EKF processes all data at the cost of being delayed. However, the mapping filter is only delayed while processing LiDAR data and will catch up with the on-line filter subsequently. Upon eventually catching up, the estimate of the mapping filter is copied to the on-line filter. This synchronization takes only less than 4 ms, thus reducing the possible delay significantly. A delay this small is again more than sufficient for path planning and control. While [5] more elegantly handles delayed measurements in a single filter the structure proposed in this paper allows for a high output rate with delays below 4 ms thanks to multi-threaded parallel processing. Inserting a delayed measurement and recalculating the subsequent updates would take up to more than 20 ms in the worst case according to the evaluation of Table I. For a desired planning and control rate of 100 Hz, this would already be too much.

Final smoothing takes about 25 seconds for the aforementioned path of about 1.8 km and more than 60,000 measurements of all sensors. The insertion of one LiDAR scan with approximately 120,000 3D points takes about 7 ms plus 32 ms for including negative information. Thus, building the whole map (Fig. 8) with about 5,000 scans takes roughly 200 seconds with integration of negative information. When operating in environments where no dynamic objects are to be expected or only mapping over short time intervals is necessary, negative information may be omitted resulting in a more than fivefold speedup.

VI. CONCLUSIONS

This paper presented an EKF based fusion framework for localization and mapping. By employing a combination of stochastic cloning and smoothing, the robust inclusion of multiple relative measurements can be achieved. Adding scan matching updates greatly improves the precision of the localization and allows for the realization of loop closure handling via GPS offset estimation. A final smoothing step of the whole path further enhances the precision for mapping.

The implementation of a synchronized parallel filter structure allows for high-rate on-line processing with low delays.

Future work includes the implementation of a parallel pipeline for scan matching to allow for processing of all scan pairs to further enhance the real-time capabilities.

ACKNOWLEDGMENT

The described research is supported by the project entitled "AKIT – Autonomy kit for near-serial-production work vehicles for networking and assisted rescue from safety hazards", which is being promoted in the course of the "Innovative rescue systems" announcement of the BMBF within the scope of the German government's "Research for civil safety" program.

REFERENCES

[1] P. Goel, S. I. Roumeliotis, and G. S. Sukhatme, "Robust localization using relative and absolute position estimates," in *Proceedings 1999 IEEE/RSJ International Conference on Intelligent Robots and Systems. Human and Environment Friendly Robots with High Intelligence and Emotional Quotients (Cat. No.99CH36289)*, vol. 2, 1999, pp. 1134–1140 vol.2.

[2] J. Wendel, *Integrierte Navigationssysteme*. München: Oldenbourg Wissenschaftsverlag, 2007.

[3] G. Welch and G. Bishop, "SCAAT: Incremental tracking with incomplete information," in *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques*, ser. SIGGRAPH '97. New York, NY, USA: ACM Press/Addison-Wesley Publishing Co., 1997, pp. 333–344.

[4] S. I. Roumeliotis and J. W. Burdick, "Stochastic cloning: a generalized framework for processing relative state measurements," in *Proceedings 2002 IEEE International Conference on Robotics and Automation (Cat. No.02CH37292)*, vol. 2, 2002, pp. 1788–1795 vol.2.

[5] S. Lynen, M. W. Achtelik, S. Weiss, M. Chli, and R. Siegwart, "A robust and modular multi-sensor fusion approach applied to mav navigation," in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Nov 2013, pp. 3923–3929.

[6] T. Bailey, J. Nieto, J. Guivant, M. Stevens, and E. Nebot, "Consistency of the EKF-SLAM algorithm," in *Proceedings of the 2006 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct 2006, pp. 3562–3568.

[7] M. Montemerlo and S. Thrun, *FastSLAM – A Scalable Method for the Simultaneous Localization and Mapping Problem in Robotics*, ser. STAR – Springer Tracts in Advanced Robotics. Berlin Heidelberg New York: Springer Verlag, 2007, vol. 27.

[8] S. Kohlbrecher, J. Meyer, O. von Stryk, and U. Klingauf, "A Flexible and Scalable SLAM System with full 3D Motion Estimation," in *Proceedings of the 2011 IEEE International Symposium on Safety, Security and Rescue Robotics (SSRR)*. IEEE, November 2011.

[9] T. Emter, "Integrated Multi-Sensor Fusion and SLAM for Mobile Robots," *Proceedings of the 2011 Joint Workshop of Fraunhofer IOSB and Institute for Anthropomatics, Vision and Fusion Laboratory*, 2012.

[10] D. Borrmann, J. Elseberg, K. Lingemann, A. Nüchter, and J. Hertzberg, "Globally consistent 3D mapping with scan matching," *Robotics and Autonomous Systems*, vol. 56, no. 2, 2008.

[11] S. Williams, V. Indelman, M. Kaess, R. Roberts, J. J. Leonard, and F. Dellaert, "Concurrent filtering and smoothing: A parallel architecture for real-time navigation and full smoothing," *Int. J. Rob. Res.*, vol. 33, no. 12, Oct. 2014.

[12] F. Moosmann and C. Stiller, "Velodyne SLAM," in *Proceedings of the 2011 IEEE Intelligent Vehicles Symposium*, Baden-Baden, Germany, June 2011, pp. 393–398.

[13] S. Rusinkiewicz and M. Levoy, "Efficient variants of the ICP algorithm," in *Proceedings Third International Conference on 3-D Digital Imaging and Modeling*, 2001, pp. 145–152.

[14] A. Segal, D. Haehnel, and S. Thrun, "Generalized-ICP." in *Robotics: Science and Systems*, vol. 2, 2009.

[15] S. Särkkä, *Bayesian Filtering and Smoothing*. Cambridge University Press, 2013.

[16] M. Magnusson, *The three-dimensional normal-distributions transform: an efficient representation for registration, surface analysis, and loop detection*. Örebro Universitet, 2009.

[17] J. E. Bresenham, "Algorithm for computer control of a digital plotter," *IBM Systems Journal*, vol. 4, no. 1, pp. 25–30, 1965.

[18] M. S. Braasch, "Multipath," in *Springer Handbook of Global Navigation Satellite Systems*, P. J. Teunissen and O. Montenbruck, Eds. Cham: Springer International Publishing, 2017, pp. 443–468.

[19] J. Kouba, F. Lahaye, and P. Tétreault, "Precise point positioning," in *Springer Handbook of Global Navigation Satellite Systems*, P. J. Teunissen and O. Montenbruck, Eds. Cham: Springer International Publishing, 2017, pp. 723–751.

[20] V. Zaliva and F. Franchetti, "Barometric and GPS altitude sensor fusion," in *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2014, pp. 7525–7529.

[21] T. Emter and J. Petereit, "Integrated Multi-Sensor Fusion for Mapping and Localization in Outdoor Environments for Mobile Robots," in *Proceedings SPIE 9121, Multisensor, Multisource Information Fusion: Architectures, Algorithms, and Applications 2014*, May 2014.