IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR),
Purdue University, West Lafayette, USA, October 2015.

Autonomous MAV Navigation in Complex GNSS-denied 3D Environments

Matthias Nieuwenhuisen, David Droeschel, Marius Beul, and Sven Behnke

***Abstract*— Micro aerial vehicles, such as multirotors, are particular well suited for the autonomous exploration, examination, and surveillance of otherwise inaccessible areas, e.g., for search and rescue missions in indoor disaster sites. Key prerequisites for the fully autonomous operation of micro aerial vehicles in restricted environments are 3D mapping, real-time pose tracking, obstacle detection, and planning of collision-free trajectories.**

In this work, we propose a complete navigation system with a multimodal sensor setup for omnidirectional environment perception. Measurements of a 3D laser scanner are aggregated in egocentric local multiresolution grid maps. Local maps are registered and merged to allocentric maps in which the MAV localizes. For autonomous navigation, we generate trajectories in a multi-layered approach: from mission planning over global and local trajectory planning to reactive obstacle avoidance. We evaluate our approach in a GNSS-denied indoor environment where multiple collision hazards require reliable omnidirectional perception and quick navigation reactions.

I. Introduction

Micro aerial vehicles (MAVs) are becoming a key factor in reducing the required time, risks, and costs for search and rescue missions, inspection tasks and aerial photography. Due to their flexibility and low cost, they constitute a superior alternative to employing heavy machinery or even risking the health of humans in many situations. Still, in most cases, a human operator pilots the MAV remotely to fulfill a specific task or the MAV is following a predefined path of GNSS waypoints in an obstacle-free altitude.

MAVs allow to quickly visit otherwise inaccessible volumes, but permanent line-of-sight to the MAV may not be maintainable. Also, passages may be narrow and surrounding environmental structures may be hard to perceive for a human operator. Hence, remotely controlling an MAV in complex 3D environments is much more demanding than controlling a ground vehicle. For example, Kruijff et al. report an overload of MAV operators while inspecting an earthquake site [1]. Narrow passages and fully closed rooms also prevent the use of global navigation satellite systems (GNSS) like GPS or GLONASS such that GNSS-based hovering or waypoint following is not an option.

In order to safely navigate in such environments, an alternative is to make MAVs autonomous, such that they can on their own—without interaction with the operator—solve well-defined sub-tasks. For example, the operator may specify a set of regions within a building and the MAV autonomously approaches all regions and collects sensor information.

This work has been supported by the German Federal Ministry for Economic Affairs and Energy (BMWi) in the Autonomics for Industry 4.0 project InventAIRy.

The authors are with the Autonomous Intelligent Systems Group, Institute for Computer Science VI, University of Bonn, Germany nieuwenh at ais.uni-bonn.de

Fig. 1: Equipped with a plurality of sensors, our MAV is capable to localize and navigate in complex indoor environments, e.g., a hall with structures hanging from the ceiling.

For the autonomous operation of MAVs, key prerequisites are localization in unmodified GNSS-denied environments, real-time obstacle detection, and planning of collision-free trajectories. Fundamental aspects in the implementation of such a system are robustness—all obstacles need to be reliably detected and mapped while avoiding false positives—and real-time onboard processing.

In this article, we present a complete integrated system consisting of a MAV with a multimodal omnidirectional sensor setup (see Fig. 1), a laser-based 3D mapping and 6D localization, and a multilayered navigation approach, tailored to the special needs of MAVs. Each layer uses and builds its own environment representation: allocentric maps for global path and mission planning and egocentric obstacle maps for local trajectory planning and reactive collision avoidance. We employ Simultaneous Localization And Mapping (SLAM) to build initial allocentric maps of the environment used for navigation and localization.

Especially in massively changed environments like damaged buildings, furniture and building debris change position and cables are hanging from the ceiling. Hence, previously built maps for planning and localization might be outdated. Thus, environment models and navigation plans have to be updated online, whenever more information becomes available during flight.

All onboard sensors have only local precision. This is reflected in the local multiresolution property of our MAV-centric obstacle map. We employ 3D local multiresolution path planning. These techniques allow for efficient map updates and frequent replanning, which makes 3D navigation

in dynamic, unpredictable environments possible.

This article extends our previous work on autonomous outdoor navigation [2]. Our main contributions here are I) robust 6D indoor localization by means of fast multilevel surfel registration and II) global navigation with a cost function—tailored to the requirements of our localization. We integrated our new components together with our fast egocentric obstacle perception and avoidance into an MAV system capable of executing fully autonomous missions in complex indoor environments. The robustness is evaluated in multiple autonomous indoor flight experiments.

II. Related Work

The use of MAVs in search-and-rescue applications is an active research topic with increasing popularity. Multiple groups use MAVs to obtain a quick overview of an area and to guide unmanned ground vehicles (UGV). Michael et al. carry an AscTec Pelican MAV on an UGV to map and inspect hardly accessible areas in a building after an earthquake [3]. Similar to our work, their MAV can operate in autonomous and semi-autonomous mode, depending on the mission requirements. However, in contrast to our work they focus on a computationally constrained MAV and omit full 3D navigation restricting it to 2.5D environments, a restriction we do not make. Luo et al. show that the guidance of UGVs by MAVs can reduce the time needed to fulfill missions [4].

Particularly important for fully autonomous operation is the ability to perceive obstacles and to avoid collisions. Obstacle avoidance is often neglected, e.g., by flying in a sufficient height when autonomously flying between waypoints.

For mobile ground robots, 3D laser scanning sensors are widely used due to their accurate distance measurements even in bad lighting conditions and their large FoV. Up to now, such 3D laser scanners are rarely used on lightweight MAVs—due to payload limitations.

We use a lightweight continuously spinning LRF that does not only allow for capturing 3D measurements without moving, but also provides omnidirectional sensing at comparably high frame rates. An MAV with a similar sensor is used by Cover et al. [5] to autonomously explore rivers using visual localization and laser-based 3D obstacle perception. In contrast to their work, we use the 3D laser scanner for both omnidirectional obstacle perception and mapping the environment in 3D.

Similar to our work, Israelsen et al. [6] present an approach to local collision avoidance that works without global localization and can aid a human operator to navigate safely in the vicinity of obstacles. Our work extends the safety layer by a deliberative planning layer based on local maps and navigation targets.

Heng et al. [7] use a multiresolution grid map to represent the surroundings of a quadrotor. A feasible plan is generated with a vector field histogram. Schmid et al. [8] autonomously navigate to user-specified waypoints in a mine. The map used for planning is created by an onboard stereo camera system. By using rapidly exploring random belief trees (RRBT), Achtelik et al. [9] plan paths that do not only avoid obstacles, but also minimize the variability of the state estimation. Recent search-based methods for obstacle-free navigation include work of MacAllister et al. [10]. They use A* search to find a feasible path in a four-dimensional grid map incorporating the asymmetric shape of the MAV. Cover et al. [5] also use a search-based method. These methods assume complete knowledge of the scene geometry—an assumption that we do not make here.

Shen et al. estimate the MAV state in combined indoor and outdoor flights by fusing different sensor modalities [11]. A 2D laser scanner is used for indoor localization. In contrast, our laser-based localization can estimate the complete 6D pose relative to a map coordinate frame.

III. System Setup and Overview

A. Platform and Sensor Setup

Our MAV is an octorotor with a co-axial arrangement of rotors (see Fig. 1). This yields a compact flying platform that is able to carry a plurality of sensors and a fast computer (Intel Core i7-3820QM 2.7 GHz). For sensor data processing and navigation planning, we employ the Robot Operating System (ROS) as middleware. For low-level velocity and attitude control, the MAV is equipped with a PIXHAWK Autopilot flight control unit. To allow for safe omnidirectional operation in challenging environments, our MAV is equipped with a multimodal sensor setup:

- Main sensor for obstacle perception is a continuously rotating 3D laser scanner. Only a small conical volume on the upper rear is occluded by the MAV core.
- Two monochrome stereo camera pairs (pointing in forward and backward direction) are used for visual odometry and obstacle perception. Equipped with fisheye lenses, they cover a large area around the MAV.
- Eight ultrasonic sensors measure distances in all directions around the MAV. Despite their limited accuracy and range of 6 m, they detect transparent obstacles, like windows.

IV. Local Perception

We construct a MAV-centric multiresolution grid map that is used to accumulate sensor measurements.

We first register newly acquired 3D scans with the so far accumulated map and then update the map with the registered 3D scan. The map is utilized by our path planning and obstacle avoidance algorithms described in subsequent sections.

A. 3D Scan Assembly

When assembling 3D scans from raw laser scans, we account for the rotation of the scanner w.r.t. the MAV and for the motion of the MAV during acquisition. Thus, scan assembling mainly consists of two steps.

First, measurements of individual scan lines are undistorted with regards to the rotation of the 2D LRF around the servo rotation axis. Here, the rotation between the acquisition

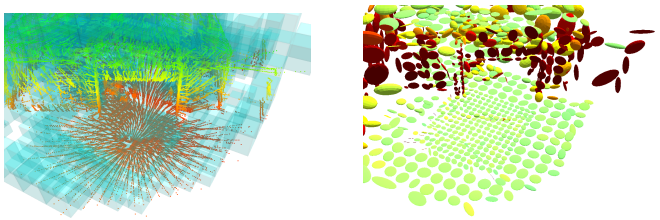

Fig. 2: Left: Grid-based local multiresolution map with a higher resolution in proximity to the sensor and a lower resolution with increasing distance. Color encodes height. Right: The surfel representation of the map (colored by surfel orientation).

of two scan lines is distributed over the measurements by using spherical linear interpolation.

Second, we compensate for the motion of the MAV during acquisition of a full 3D scan. To this end, we incorporate a visual odometry estimate from the two stereo cameras. Here, a keyframe-based bundle adjustment is performed [12] on the synchronized images with 18 Hz update rate. Since the update rate of the 2D LRF is 40 Hz, we linearly interpolate between the estimates of the visual odometry. The 6D motion estimate is used to assemble the individual 2D scan lines of each half rotation to a 3D scan.

B. Local Multiresolution Map

We use a hybrid local multiresolution map that represents both occupancy information and the individual distance measurements. The most recent measurements are stored in ring buffers within grid cells that increase in size with distance from the robot center. Thus, we use a high resolution in the close proximity to the sensor and a lower resolution far away from our robot, which correlates with the sensor characteristics in measurement accuracy and density. Compared to uniform grid-based maps, multiresolution leads to the use of fewer grid cells, without losing relevant information and consequently results in lower computational costs. Fig. 2 shows an example of our local multiresolution grid-based map.

We aim for efficient map management for translation and rotation. To this end, individual grid cells are stored in a ring buffer to allow shifting of elements in constant time. We interlace multiple ring buffers to obtain a map with three dimensions. The length of the ring buffers depends on the resolution and the size of the map. In case of a translation of the MAV, the ring buffers are shifted whenever necessary to maintain the egocentric property of the map. For sub-cell-length translations, the translational parts are accumulated and shifted if they exceed the length of a cell.

C. Registration Approach

We register each newly acquired 3D scan with the local multiresolution map of the environment with our surfel-based method [13]. Instead of considering each point individually, we represent the 3D scan as local multiresolution grid and match surfels. In each cell of the grid, we maintain one surfel that summarizes the individual 3D points that lie within the cell (cf. Fig. 2). A surfel is defined by the sample mean and the sample covariance of these points. We align a newly acquired scan (scene) and the local multiresolution map (model) by finding a rigid 6 degree-of-freedom (DoF) transformation $T(\theta)$ that best aligns the scene surfels to the model surfels. By summarizing measurements in surfels, and therefore considering several orders of magnitudes less elements for registration, we gain efficiency. When matching surfels we choose the finest common resolution available between both maps to achieve accuracy. We cope with the sparsity of the 3D scans obtained from our LRF through probabilistic assignments of surfels during the registration process.

V. Allocentric Mapping and Localization

For fast estimation of the MAV motion, we incorporate IMU and visual odometry measurements into velocity and pose estimates. While these estimates allow us to control the MAV and to track its pose over a short period of time, they are prone to drift and thus are not suitable for localization on the time scale of a mission. Furthermore, they do not provide a fixed allocentric frame for the definition of mission-relevant poses independent from the MAV. Thus, we build an allocentric map by means of laser-based SLAM before mission execution and employ laser-based pose tracking w.r.t. this map during autonomous operation.

A. Mapping

This allocentric map is build by aligning multiple local multiresolution maps, acquired from different view poses [14]. We model the different view poses as nodes in a graph that are connected by edges. A node consists of the local multiresolution map from the corresponding view pose. Each edge in the graph models a spatial constraint between two nodes.

After adding a new 3D scan to the local multiresolution map as described in Sec. IV-C, the local map is registered towards the previous node in the graph using the multiresolution surfel registration with probabilistic assignments [13].

The resulting pose graph is efficiently optimized by using the g^{2}o framework by Kuemmerle et al. [15]. After the MAV has traversed the environment, the allocentric map is build from the optimized pose graph by merging all local surfel maps. Here, we use surfels with uniform resolution. Fig. 3 shows an example map acquired from a combined indoor and outdoor flight.

B. Pose Tracking

When executing a mission, the MAV traverses a set of goal poses w.r.t. the coordinate frame defined by our allocentric map. Since the laser scanner acquires complete 3D scans with a relatively low frame rate, we incorporate the egomotion estimate from the visual odometry and measurements from the IMU to track the pose of the MAV. The egomotion estimate is used as a prior for the motion between two consecutive 3D scans. In detail, we track the pose hypothesis by alternating the prediction of the MAV movement given the

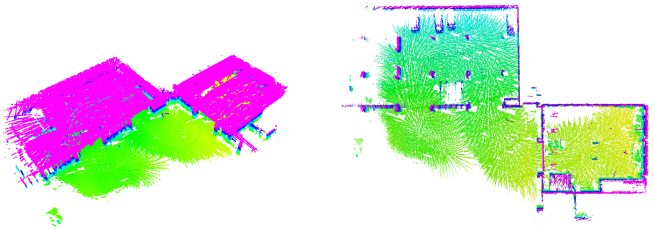

Fig. 3: Allocentric map from a combined indoor/outdoor flight after pose graph optimization. Left: Perspective view. Right: Top-view with removed ceiling.

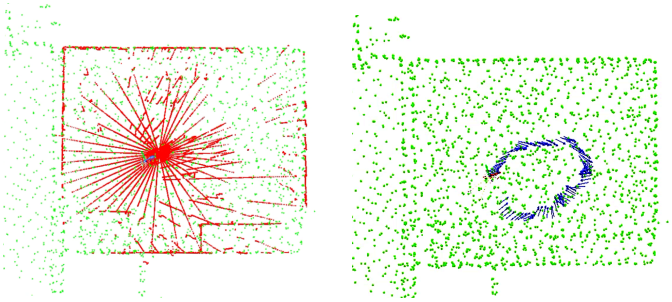

Fig. 4: Undistorted local 3D laser scans (red) are matched to the allocentric SLAM map (green) to localize the robot in the allocentric frame. Here, a downsampled version of the allocentric map is shown. Left: The estimated robot pose is depicted by the axes. Right: Blue arrows depict the estimated MAV trajectory.

filter result and alignment of the current local multiresolution map towards the allocentric map of the environment. To align the current local map with the allocentric map, we also use the surfel-based registration [13]. The allocentric localization is triggered after acquiring a 3D scan and adding it to the local multiresolution map. We update the allocentric robot pose with the resulting registration transform. To achieve real-time performance of the localization module, we track only one pose hypothesis. We assume that the initial pose of the MAV is known, either by starting from a predefined pose or by means of manually setting the pose. Fig. 4 shows the registration of a 3D scan to the map and an estimated 6D trajectory.

The resulting robot pose estimate from the allocentric localization is used as a measurement update in a lower-level state estimation filter. We propagate this allocentric pose over time with visual odometry and IMU to obtain allocentrically consistent pose and velocity estimates at a sufficiently high rate for planning and control.

VI. Navigation Planning

To operate MAVs in indoor environments for search and inspection missions, safe navigation in the vicinity of obstacles is key. Compared to outdoor missions, the free space is restricted and keeping a large safety margin to obstacles is not an option. Hence, only quick reactions given the observed vehicle and environment state ensure the successful and safe mission completion. To allow for both deliberative planning and quick reactions, we employ hierarchical navigation: from mission planning to low-level motion control. These tasks require different abstractions of the environment.

To plan an observation mission, we need a coarse model of the environment; to plan collision-free paths, we need a finer and up-to-date consistent geometric model; and to avoid collisions, we need a non-aggregated local representation of the close vicinity of the MAV. The planned actions also have different granularity, which is represented by the planning frequency, from once per mission to multiple times per second. The higher-layer planners set goals for the lower-level planners which produce more concrete action sequences based on more local and up-to date environment representations.

A. Mission Planning

The topmost layers are a mission planner and a global path planner. Both use a similar representation of the environment that is built by means of SLAM by the MAV beforehand, as described in the previous section.

Input to the mission planner is a set of view poses defined by the user. The mission planner employs a global path planner on a coarse uniform grid map to determine the approximate costs between every pair of mission goals.

The result of mission planning is a flight plan composed of an ordered list of 4D waypoints.

B. Global Path Planning

The next layer in the planning hierarchy is a global path planner. This layer plans globally consistent plans, based on I) the (updated) environment model, discretized to grid cells with $0.5\,\mathrm{m}$ edge length, II) the current pose estimate of the MAV, and III) the next mission waypoint, including 3D position, yaw orientation, and required accuracies. Planning frequency is $0.2\,\mathrm{Hz}$ and we use the A* algorithm to find cost-optimal paths.

In our application domain, most obstacles not represented in the allocentric map can be surrounded locally, without the need for global replanning. Hence, it is sufficient to replan globally on a more long-term time scale to keep the local deviations of the planner synchronized to the global plan and to avoid the MAV to get stuck in a local minimum that the local planner cannot solve due to its restricted view of the environment.

For application-specific SLAM maps, it is often not necessary to cover the whole reachable environment, but only the parts that are relevant for the mission execution. In particular, they cannot cover the complete space outside the buildings. Our laser-based localization needs to perceive sufficient structure to work robustly. Hence, the MAV should not fly in completely unmapped or free space, e.g., at a height where the ground is no longer observed by the LRF. To ensure robust localization even in partial maps and unbound environments, we employ an approach inspired by coastal navigation [16]. With increasing distance to obstacles, our obstacle cost function $h_c(d)$ decreases until distance D_o. Starting at a distance D_{p1}, the perception cost function h_p

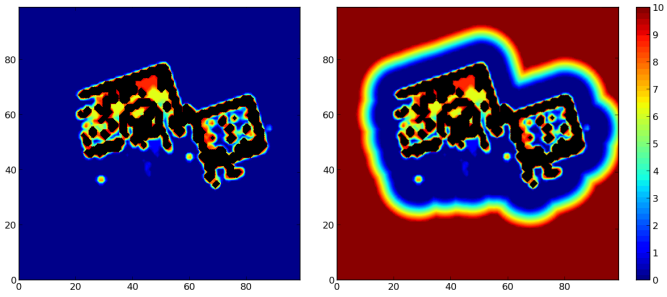

Fig. 5: Cut through cost map some meters above ground. For robust LRF-based localization, laser scans have to contain sufficient structure. We modify the traversal costs to not only avoid obstacles, but also keep them within sensor range. Left: Obstacle avoidance only. Right: Obstacle avoidance and robust localization.

increases up to a maximum at D_{p2} to keep the obstacles in the observable range of the MAV. Our cost model $h(d)$ is:

$$h_c(d) = \begin{cases} \infty & \text{if } d \leq D_s, \\ h_{max}\frac{1-d+D_s}{D_o-D_s} & \text{if } D_s < d < D_o, \\ 0 & \text{otherwise;} \end{cases}$$

$$h_p(d) = \begin{cases} 0 & \text{if } d \geq D_{p1}, \\ h_{max}\frac{d-D_{p1}}{D_{p2}-D_{p1}} & \text{if } D_{p1} < d < D_{p2}, \\ h_{max} & \text{otherwise;} \end{cases}$$

$$h(d) = w_1 \cdot h_c(d) + w_2 \cdot h_p(d),$$

with equal weights $w_1 = w_2$. D_s is the safety distance around obstacles the robot should never enter (at least the robot radius). Fig. 5 illustrates the resulting traversal costs with and without our approach.

C. Local Multiresolution Path Planning

On the local path planning layer, we employ a 3D local multiresolution path planner [17]. This layer plans based on the allocentric path from the global path planner and the 3D local multiresolution map. It refines the global path according to the actual situation and a finer trajectory is fed to the potential field-based reactive obstacle avoidance layer on the next level.

We embed an undirected graph into the multiresolution grid and perform A* search from the origin of the MAV-centered grid to the goal. The traversal costs h_c of a cell with distance d to the robot are modeled as

$$h_c(d) = h_{max}\frac{1-d-r_C}{2*r_C}$$

for $d \in [r_C, 3*r_C]$, where r_C is the perceived obstacle with a distance-dependent radius that models the uncertainty of farther away perceptions. Closer and farther away obstacles are modeled with cost h_{max} and 0, respectively.

The allocentric path is cost-optimal with respect to the allocentric map. Hence, the path costs of the allocentric path are a lower bound to path costs for refined plans, based on newly acquired sensor information—mostly dynamic and static previously unknown obstacles—and a local path deviating from the global plan cannot be shorter in terms of path costs. Locally shorter plans on lower layers with a local view on the map may yield globally suboptimal paths. We address this by coupling the local planner to the solution of the allocentric path planner by a cost term h_a, which is the shortest distance between a grid cell and any segment of the allocentric plan. The total cost h for traversing a grid cell is $h = w_1 \cdot h_c(d) + w_2 \cdot h_a$.

Local multiresolution path planning leverages efficiency. Since parts of the plan that are farther away from the MAV are more likely to change, e.g., due to newly acquired sensor data, it is reasonable to spend more effort into a finer plan in the close vicinity of the robot. Overall, our approach reduces the planning time and makes frequent replanning feasible.

D. Local Obstacle Avoidance

On the next lower layer, we employ a fast reactive collision avoidance module based on artificial potential fields as a safety measure reacting directly on the available sensor inputs. The robot-centered local multiresolution occupancy grid, the current motion state, and a target velocity, serve as input to our algorithm. The obstacle map induces repulsive forces on the MAV pushing it away from obstacles. Analysis of the MAV motion model gives us a maximum stopping distance of 0.8 m in 0.5 s under ideal conditions. Hence, we can safely operate in the vicinity of obstacles. For more details on the potential field obstacle avoidance layer, please refer to our prior work [17].

VII. Evaluation

We evaluated the integration of our components into one working MAV mapping and surveillance system. Here, we describe one exemplary indoor surveillance session of our MAV. The main goal of this experiment was to autonomously navigate to certain predefined waypoints in a hall, employing solely means of localization that are available in both indoor and outdoor environments. During flight, the MAVs mission was to detect visual features (AprilTags [18]), representing interesting locations near the trajectory. Local navigation and control were performed using visual odometry and allocentric localization was performed employing our laser-based approach.

First, we navigated the MAV manually through a hall, a garage, and an outdoor part connecting these two buildings. Subsequently, we built a map from the data collected during this flight (Fig. 3). The map is used for localization and we derived an OctoMap for mission and path planning. In applications where such an initial flight is not feasible, building construction plans could be used instead.

Second, we defined a mission with six observation poses—plus a return pose 2 m above the start pose—in the smaller building part, a decommissioned car service station with multiple collision hazards. Our mission planner plans paths between every pair of mission poses and determines the best visiting order. After takeoff, the global planner begins to continuously plan paths to the next mission-relevant pose. The local obstacle avoidance keeps the MAV successfully

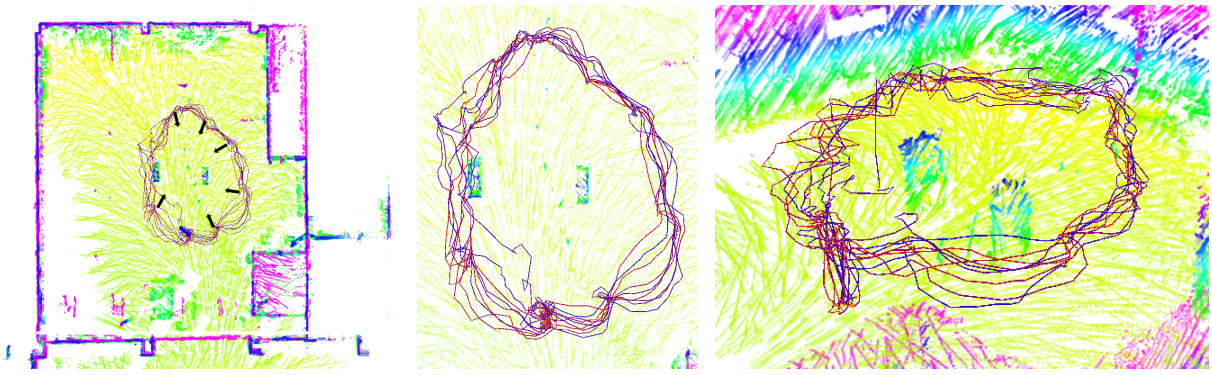

Fig. 6: We conducted indoor flight experiments in a $20 \times 16 \times 6\,\mathrm{m}$ hall. The trajectories show the laser-based localization estimates of 10 consecutive autonomous flights. From left to right: map top-view, closeup of the trajectories, perspective view. Observation poses are depicted by black arrows, the blue arrow shows the return pose. The map color encodes height.

Fig. 7: The MAV is repelled from a small pipe structure (red circled) hanging from the ceiling close to the MAV's view pose. The right figure shows how the obstacle is perceived by the MAV. Red lines depict the artificial repelling forces of the reactive collision avoidance.

away from obstacles like hanging cables and debris, lying around in the hall. In these experiments, we planned allocentric paths in a grid with a cell size of 0.5 m. An excerpt from the map, the inspection poses, and the traversed trajectories of ten missions are shown in Fig. 6. The MAV successfully reached all poses in the experiments without colliding with an obstacle and was localized all the time. In experiments without running localization module, the MAV still was able to avoid dynamic and static obstacles with the local collision avoidance layer. This layer solely relies on egocentric velocity estimates, e.g., from an integration of visual odometry, accelerometers, and gyroscopes.

Fig. 7 shows a situation from one of the runs where the MAV is flying towards a mission view pose. Close to the view pose, a long thin structure is hanging from the ceiling. The structure is perceived with the onboard sensors and avoided by means of artificial repelling forces. In our test setup, many smaller and larger structures obstruct the free space. In all cases, the MAV deviated from the direct path or moved away from a hover position to avoid a collision. Fig. 8 shows the magnitudes of repelling forces while fulfilling a planned mission. The shown positions are from the ten autonomous flights shown in Fig. 6 plus one additional mission with observation poses closer to some obstacles. Samples for illustration are taken every 500 ms. A video of our mission execution experiments is available as attachment to this work[1]. Further collision avoidance experiments are presented in a video on our website[2].

During the execution of the surveillance mission, the MAV detects visual features. As seen in Fig. 8, we use AprilTags with $26\,\mathrm{cm}$ side length as a surrogate for interesting locations near the trajectory. Here, they represent, e.g., victims in an earthquake-scenario, pipes leaking poisonous liquids, or burning debris. Surrogating visual features is a common technique used to evaluate mapping systems, e.g., in RoboCup Rescue[3].

Fig. 9 shows a map of the detected AprilTags during ten flights. It can be seen that the clustering of all detections is within $0.5\,\mathrm{m}$ and that most tags are detected many times. So, even if, e.g., smoke would prevent the perception from one side, the tag could nevertheless be perceived. The detection and mapping of AprilTags is detailed in [19].

VIII. Conclusions

In this work, we presented an integrated system to autonomously operate MAVs safely in the vicinity of obstacles. We approached this challenge by employing local multiresolution mapping and planning techniques that facilitate frequent updates and replanning.
We showed that by incorporating multimodal sensor information, we are able to detect and avoid diverse obstacles. Pose estimation based on camera and laser data enables robust motion control in GNSS-denied environments.

Laser-range measurements are aggregated by registering sparse 3D scans with a local multiresolution surfel map. Modeling measurement distributions within voxels by surface elements allows for efficient and accurate registration of 3D scans with the local map. The incrementally built local dense 3D maps of nearby key poses are registered globally by graph optimization. This yields a globally consistent dense 3D map of the environment. We demonstrate accuracy and efficiency of our approach by showing consistent allocentric 3D maps, recorded by our MAV during flight.

We demonstrated that multilayered navigation planning results in a high capability to cope with dynamically changing environments and perpetually new obstacle perceptions. A

[1] www.ais.uni-bonn.de/InventAIRy/videos/ssrr-indoor.mp4

[2] www.ais.uni-bonn.de/MoD

[3] www.robocuprescue.org

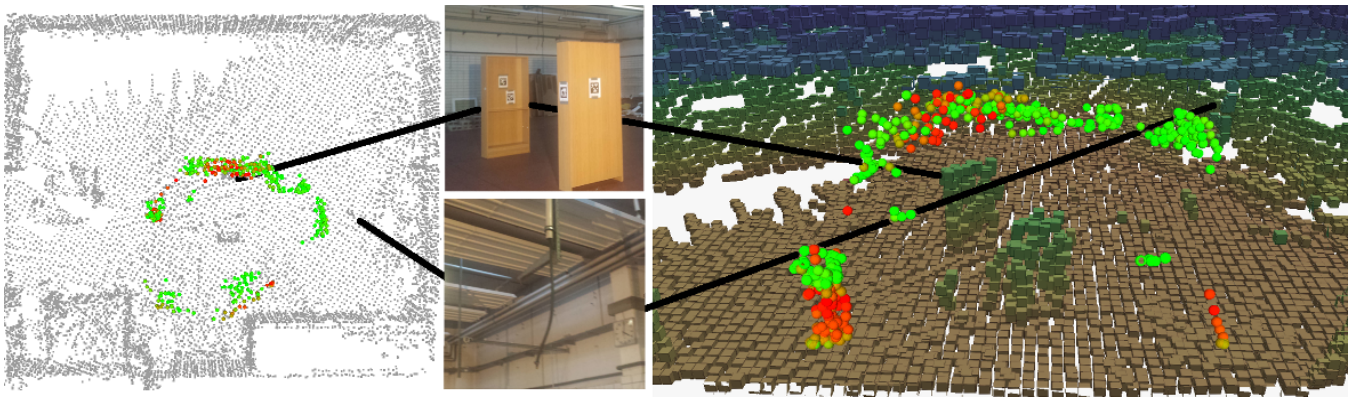

Fig. 8: Artificial forces from the reactive collision avoidance. Shown is the strength of repelling forces pushing the MAV away from obstacles at samples along the trajectories of ten autonomous flights. Green: small magnitude / Red: large magnitude. The photos show some obstacles and their approximate position in the maps.

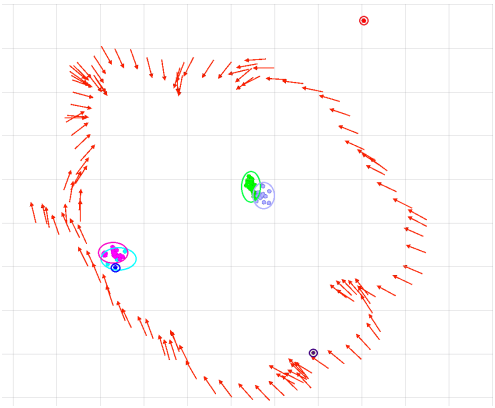

Fig. 9: Map of detected AprilTags of ten consecutive flights. Seven different tags are perceived and marked in magenta, cyan, blue, green, gray, violet and red. A sample trajectory is shown in red. The grid size is $1 \times 1\,\mathrm{m}$.

reactive collision avoidance layer accounts for fast MAV and environment dynamics and refines higher-level mission plans based on onboard sensing and a priori information.

We showed the system robustness in multiple indoor experiments where the only manual interactions were the starting and landing phases. Thus, the system is able inspect areas for the detection of threats or search for victims in a fully autonomous mission. This reduces the workload of human operators in disaster situations.

References

[1] G.-J. M. Kruijff, I. Kruijff-Korbayová, S. Keshavdas, B. Larochelle, M. Janíček, F. Colas, M. Liu, F. Pomerleau, R. Siegwart, M. A. Neerincx, R. Looije, N. J. J. M. Smets, T. Mioch, J. van Diggelen, F. Pirri, M. Gianni, F. Ferri, M. Menna, R. Worst, T. Linder, V. Tretyakov, H. Surmann, T. Svoboda, M. Reinštein, K. Zimmermann, T. Petříček, and V. Hlaváč, "Designing, developing, and deploying systems to support human–robot teams in disaster response," *Advanced Robotics, special issue on Disaster Response Robotics*, vol. 28, no. 23, pp. 1547–1570, December 2014.

[2] D. Droeschel, M. Nieuwenhuisen, M. Beul, D. Holz, J. Stückler, and S. Behnke, "Multi-layered mapping and navigation for autonomous micro aerial vehicles," *J. of Field Robotics*, 2015, published online.

[3] N. Michael, S. Shen, K. Mohta, V. Kumar, K. Nagatani, Y. Okada, S. Kiribayashi, K. Otake, K. Yoshida, K. Ohno, E. Takeuchi, and S. Tadokoro, "Collaborative mapping of an earthquake-damaged building via ground and aerial robots," in *Int. Conf. on Field and Service Robotics (FSR)*, 2012.

[4] C. Luo, A. P. Espinosa, D. Pranantha, and A. D. Gloria, "Multi-robot search and rescue team," in *Proc. of the IEEE Int. Symposium on Safety, Security and Rescue Robotics (SSRR)*, 2011.

[5] H. Cover, S. Choudhury, S. Scherer, and S. Singh, "Sparse tangential network (SPARTAN): Motion planning for micro aerial vehicles," in *Proc. of the Int. Conf. on Robotics and Automation (ICRA)*, 2013.

[6] J. Israelsen, M. Beall, D. Bareiss, D. Stuart, E. Keeney, and J. van den Berg, "Automatic collision avoidance for manually tele-operated unmanned aerial vehicles," in *Proc. of the Int. Conf. on Robotics and Automation (ICRA)*, 2014.

[7] L. Heng, D. Honegger, G. H. Lee, L. Meier, P. Tanskanen, F. Fraundorfer, and M. Pollefeys, "Autonomous visual mapping and exploration with a micro aerial vehicle," *J. of Field Robotics*, vol. 31, no. 4, pp. 654–675, 2014.

[8] K. Schmid, P. Lutz, T. Tomic, E. Mair, and H. Hirschmüller, "Autonomous vision-based micro air vehicle for indoor and outdoor navigation," *J. of Field Robotics*, vol. 31, no. 4, pp. 537–570, 2014.

[9] M. W. Achtelik, S. Lynen, S. Weiss, M. Chli, and R. Siegwart, "Motion- and uncertainty-aware path planning for micro aerial vehicles," *J. of Field Robotics*, vol. 31, no. 4, pp. 676–698, 2014.

[10] B. MacAllister, J. Butzke, A. Kushleyev, H. Pandey, and M. Likhachev, "Path planning for non-circular micro aerial vehicles in constrained environments," in *Proc. of the Int. Conf. on Robotics and Automation (ICRA)*, 2013.

[11] S. Shen, Y. Mulgaonkar, N. Michael, and V. Kumar, "Multi-sensor fusion for robust autonomous flight in indoor and outdoor environments with a rotorcraft mav," in *Proc. of the Int. Conf. on Robotics and Automation (ICRA)*, 2014.

[12] J. Schneider, T. Läbe, and W. Förstner, "Incremental real-time bundle adjustment for multi-camera systems with points at infinity," in *Int. Arch. Photogramm. Remote Sens. Spatial Inf. Sci. (ISPRS)*, vol. XL-1/W2, 2013.

[13] D. Droeschel, J. Stückler, and S. Behnke, "Local multi-resolution representation for 6D motion estimation and mapping with a continuously rotating 3D laser scanner," in *Proc. of the Int. Conf. on Robotics and Automation (ICRA)*, 2014.

[14] ——, "Local multi-resolution surfel grids for mav motion estimation and 3D mapping," in *Proc. of the Int. Conf. on Intelligent Autonomous Systems (IAS)*, 2014.

[15] R. Kuemmerle, G. Grisetti, H. Strasdat, K. Konolige, and W. Burgard, "g^2o: A general framework for graph optimization," in *Proc. of the Int. Conf. on Robotics and Automation (ICRA)*, 2011.

[16] N. Roy, W. Burgard, D. Fox, and S. Thrun, "Coastal navigation-mobile robot navigation with uncertainty in dynamic environments," in *Proc. of the Int. Conf. on Robotics and Automation (ICRA)*, 1999.

[17] M. Nieuwenhuisen and S. Behnke, "Layered mission and path planning for MAV navigation with partial environment knowledge," in *Proc. of the Int. Conf. on Intelligent Autonomous Systems (IAS)*, 2014.

[18] E. Olson, "AprilTag: A robust and flexible visual fiducial system," in *Proc. of the Int. Conf. on Robotics and Automation (ICRA)*, 2011.

[19] M. Beul, N. Krombach, Y. Zhong, D. Droeschel, M. Nieuwenhuisen, and S. Behnke, "A high-performance MAV for autonomous navigation in complex 3D environments," in *Proc. of the Int. Conf. on Unmanned Aircraft Systems (ICUAS)*, 2015.