In Proceedings of the 6th European Conference on Mobile Robots (ECMR), Barcelona, Spain, 2013

Multimodal Obstacle Detection and Collision Avoidance for Micro Aerial Vehicles

Matthias Nieuwenhuisen[1], David Droeschel[1], Johannes Schneider[2], Dirk Holz[1], Thomas Läbe[2], and Sven Behnke[1]

***Abstract*— Reliably perceiving obstacles and avoiding collisions is key for the fully autonomous application of micro aerial vehicles (MAVs). Limiting factors for increasing autonomy and complexity of MAVs are limited onboard sensing and limited onboard processing power. In this paper, we propose a complete system with a multimodal sensor setup for omnidirectional obstacle perception. We developed a lightweight 3D laser scanner and visual obstacle detection using wide-angle stereo cameras. Detected obstacles are aggregated in egocentric grid maps. We implemented a fast reactive collision avoidance approach for safe operation in the vicinity of structures like buildings or vegetation.**

I. Introduction

In the context of a larger project on three-dimensional semantic mapping of inaccessible areas and objects, we aim at developing a micro aerial vehicle (MAV) that is able to autonomously navigate in suburban areas and especially in the (close) vicinity of buildings, vegetation and other possibly dynamic objects. In particular, we focus on fast and reliable perception of (even small) obstacles in the vicinity of the MAV. In this paper, we describe the hardware design of our platform including the sensor setup and the applied methods for robust obstacle detection, as well as for planning and controlling the MAV's motion in order to reach goal poses while reliably avoiding collisions.

MAVs such as quadrotors have attracted much attention in the field of aerial robotics due to low cost and ease of control. Their size and weight limitations, however, pose a problem in designing sensory systems for these robots. Most of today's MAVs are equipped with ultra sound sensors and cameras due to their minimal size and weight. In contrast, authors in [1], [2], [3], [4] have equipped their platforms with 2D laser range finders (LRF) that are used for navigation.

2D scanners are restricted to a measurement plane. In the general field of robotics, however, three-dimensional (3D) laser scanning sensors are widely accepted for mobile robots due to their accurate distance measurements even in bad lighting conditions and their large field-of-view (FoV). For instance, autonomous cars usually perceive obstacles by means of a rotating laser scanner with a 360° horizontal FoV, allowing to detect obstacles in every direction. Up to now, such 3D laser scanners are too heavy to be used on MAVs.

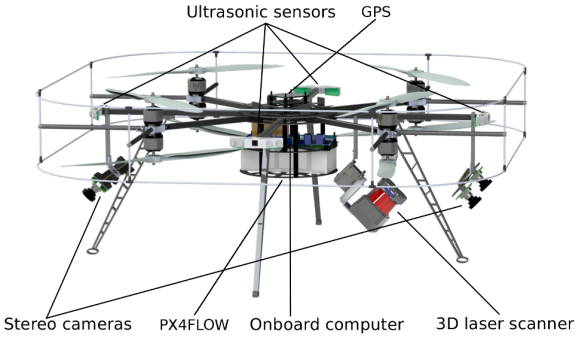

Fig. 1: Our MAV is equipped with eight co-axial rotors and a plurality of sensors, including a continuously rotating 3D laser scanner and two stereo camera pairs.

We have designed a continuously rotating laser scanner that is minimalistic in terms of size and weight and is particularly well suited for obstacle perception and localization on MAVs, allowing to perceive the environment in all directions. Other means of obstacle perception on our MAV are ultrasonic sensors and cameras. Fig. 1 shows the sensors mounted on our MAV to perceive obstacles in its environment and to estimate the ego-motion of the robot.

II. Related Work

The application of MAVs in recent robotics research varies especially in the level of autonomy ranging from basic hovering and position holding [5] over trajectory tracking and waypoint navigation [6] to fully autonomous navigation [7]. Limiting factors for increasing autonomy and complexity of truly autonomous systems (without external sensing and control) are limited onboard sensing and limited onboard processing power.

Particularly important for fully autonomous operation is the ability to perceive obstacles and avoid collisions. Most autonomous MAVs, however, cannot adequately perceive their surroundings and autonomously avoid collisions. Instead, collision avoidance is often restricted to the two-dimensional measurement plane of laser range finders [7] or the limited field of view of (forward-facing) cameras—or generally avoided, e.g., by flying in a certain height when autonomously flying between waypoints. Tomić *et al.* present an autonomous aerial vehicle that perceives its environments using a stereo camera pair mounted forwards and a 2D laser

This work has been supported as part of the research group FOR 1505 of German Research Foundation (DFG).

[1]Autonomous Intelligent Systems Group, Computer Science Institute VI, University of Bonn, 53113 Bonn, Germany {nieuwenh, droeschel, holz} at ais.uni-bonn.de, behnke at cs.uni-bonn.de

[2]Photogrammetry, Institute of Geodesy and Geoinformation, University of Bonn, 53115 Bonn, Germany josch at uni-bonn.de, laebe at ipb.uni-bonn.de

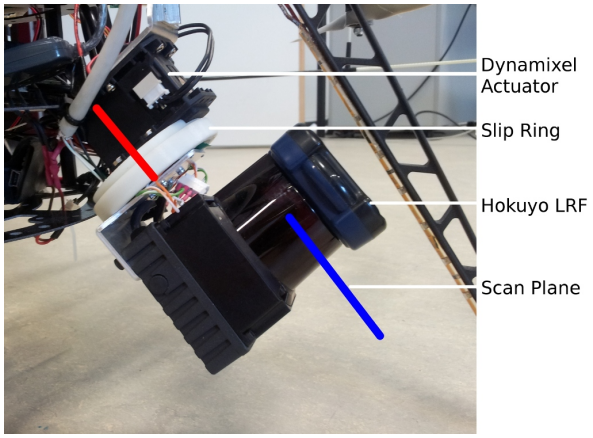

Fig. 2: Our continuously rotating 3D laser scanner is based on a Hokuyo 2D LRF mounted on a bearing. This allows to rotate the measurement plane (blue line) 360° around the red axis to acquire complete 3D scans of the environment.

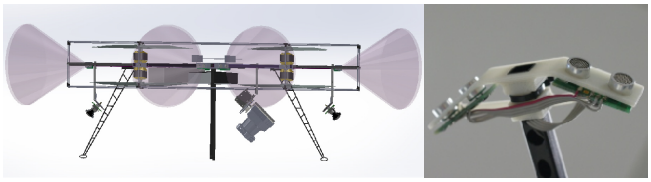

(a) CAD model with measurement cones (b) Photo

Fig. 3: Setup and mounting of ultrasonic sensors.

range scanner mounted horizontally [1]. Still, their perceptual field is limited to the apex angle of the stereo camera pair (facing forwards), and the 2D measurement plane of the scanner when flying sideways. They do not perceive obstacles outside of this region or behind the vehicle. We aim at perceiving as much of the surroundings as possible in order to obtain almost omnidirectional obstacle detection.

III. Sensor Setup

To perceive obstacles reliably, we incorporate different sensor modalities into our system. Our main sensor consists of a Hokuyo UTM-30LX-EW 2D LRF that is mounted on a bearing plate which is connected to a Dynamixel MX-28 servo actuator. The servo actuator continuously rotates the 2D LRF to gain a three-dimensional FoV. To enlarge the FoV and reduce occlusions, the LRF is slightly shifted and twisted away from the axis of rotation (see Fig. 2). The 2D LRF is connected to the system by a slip ring, allowing for seamless rotation of the sensor. The whole setup is pitched downward by 45° which allows to maximize the FoV and to minimize the blind spot of the sensor.

A single 2D scan of the LRF consists of 1080 distance measurements (270° apex angle and 0.25° angular resolution) and is called a *scan line*. The LRF is rotated with one rotation per second, resulting in 40 scan lines and 43200 distance measurements per full rotation. Since a half rotation leads to a full 3D scan of the environment, we can acquire 3D scans with 21600 points at a rate of 2 Hz.

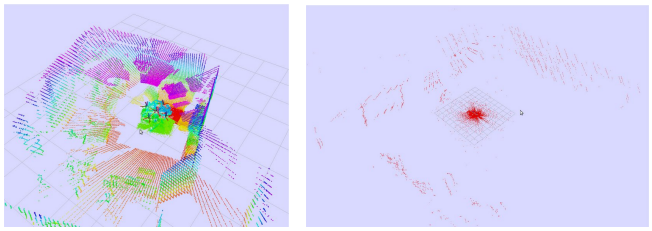

(a) (b)

Fig. 4: 3D scans acquired with our continuously rotating laser scanner in an indoor environment (a) and in an outdoor environment during flight (b).

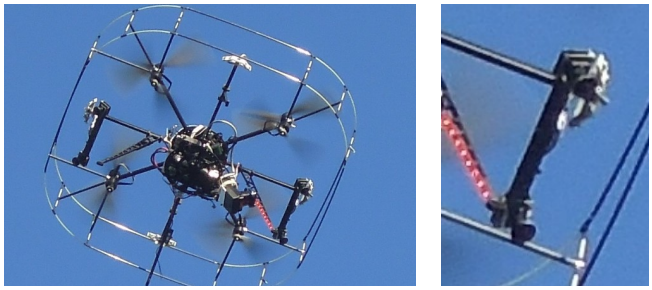

Fig. 5: Illustration of the MAV, one stereo pair is looking forward and one backwards giving an omnidirectional field of view.

The rotation angle of the sensor is measured by the actuator's encoder. With that, we can calculate a 3D point cloud from the distance measurements. Resulting point clouds in indoor and outdoor environments are shown in Fig. 4.

For visual obstacle detection, we have mounted four cameras with fisheye lenses with a viewing angle up to 185° on the MAV to generate two stereo pairs. One stereo pair is looking ahead and one is looking backwards, providing an omnidirectional field of view around and below the robot, as shown in Fig. 5. The baseline of the stereo pairs is 20 cm. The monochromatic cameras capture four image sequences with a frame rate of 14 Hz in a synchronized way. A master camera triggers the image acquisition.

As neither the laser point cloud nor the visual obstacles are dense, our MAV is equipped with eight ultrasonic sensors covering the near space around it (see Fig. 3). These sensors detect smaller obstacles in the vicinity of the robot, e.g. wires. They also measure transparent objects, like windows, which are problematic for visual sensors.

IV. Visual Obstacles

Visual obstacle detection is based on interest points, which are tracked in the video streams of the cameras using the OpenCV implementation of the KLT tracker and matched across the cameras, if possible. Interest points are corners in the gradient image with a large smallest eigenvalue of the structure tensor, see [8]. Tracking is performed by using the iterative Lucas-Kanade method with pyramids according to [9]. Besides obstacle detection, the tracked features are used for visual odometry—a topic we will not discuss in this paper. For visual obstacle detection, we use the matched feature points to determine the coordinates of the observed

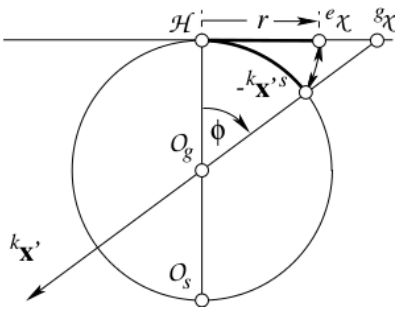

Fig. 6: Relation between sensor point, viewing direction and viewing ray.

scene points in the camera frame at every time of exposure via stereo triangulation. The mutual orientations of the cameras in a stereo pair are determined in advance according to [10].

Each feature point is converted into a ray direction pointing to the observed scene point in the individual camera frame system. For this, we model the fisheye lens with the equidistant model described in [11]—allowing for ray directions with an intersection angle equal or higher than $90°$ to the viewing direction. The interior orientation of each camera is determined in advance by camera calibration according to [12] using Chebyshev polynomials. Using the equidistant projection and applying all corrections to the feature points, we obtain image points $^e\chi$ lying closer to the principal point $\mathcal{H}$ than the gnomonic projections $^g\chi$ of the observed scene points, see Fig. 6. The ray direction $^k\boldsymbol{x}'^s$ pointing to the unknown scene points in the camera frame system can be derived from $^e\chi$ by using the normalized radial distance $r' = |^e\boldsymbol{x}|$ growing with the angle ϕ between the viewing direction and the camera ray.

To match feature points in the overlapping images of a stereo camera pair, we determine the correlation coefficients between the local 7×7 image patches at the feature points in the left and right images. Using the known relative orientation between the cameras within a stereo pair, we can reduce the amount of possible candidates to feature points lying close to the corresponding epipolar lines, see Fig.7. We assume feature points with the highest correlation coefficient ρ_1 to match, if ρ_1 is above an absolute threshold, e. g. 0.8, and—if there is more than one candidate close to the epipolar line—the closest-to-second-closest-ratio $r = \rho_2/\rho_1$ with the second highest correlation coefficient ρ_2 is lower than an absolute threshold, e. g. 0.7.

For fast obstacle detection, we determine an unknown scene point by directly intersecting the corresponding camera rays $^kx'^s$ using the known mutual orientations between the cameras within a stereo pair. This procedure leads in some rare cases to wrong matches, which can only be detected with a third observing ray from another pose.

V. Sensor Data Acquisition and Processing Pipeline

In this section, we discuss our processing pipeline—from the acquisition of sensor measurements to updating the egocentric obstacle map.

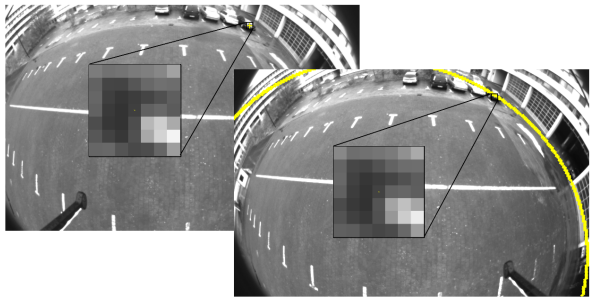

Fig. 7: Example stereo image pair. The extracted feature point in the left image on the rightmost car has the illustrated epipolar line in the right image. The matching point in the right image lies on the indicated yellow line and the corresponding local image patches show a high correlation.

A. Self Filtering and Removing Erroneous Measurements

Because of the large FoV of the sensors, a considerable amount of points is either measured directly on the robot itself, or caused by occlusion effects. We filter out such measurements by applying a simplified robot model for estimating which measurements coincide with the robot's body parts. Referring to Fig. 8, we distinguish between measured points on the aerial vehicle and measured points belonging to obstacles in the robot's vicinity.

B. Laser-based Height Estimation

In order to obtain an accurate height estimate, we first compute the set of points below the robot and then find the most dominant (horizontal) plane for these points. We define a downwards pointing frustum under the robot. For the points within the frustum, we apply an approach based on the M-Estimator sample consensus (MSAC) [13]. After finding the most dominant plane model, we determine the set of inliers (from the complete laser scan) for the found plane model, and then refine the model by fitting a plane through all inliers. We use the distance of the robot to the estimated ground plane as a height estimate within our state estimation approach. The runtimes for scan acquisition, filtering and ground plane estimation as described above lie in the range of milliseconds.

C. Obstacle Map

The distance measurements of the sensors are accumulated in a 3D grid-based map which is centered around the robot and oriented independent of the vehicle's orientation. We maintain the orientation between the map and the MAV and use it to rotate measurements when accessing the map. For each sensor measurement, the individual cell of the map is marked as occupied. We move occupied cells according to the estimated ego-motion of the robot.

VI. Local Obstacle Avoidance

Our concept for the navigation of the MAV is based on a multi-layer approach. Between low-level control and high-level planning layers, we employ a fast reactive collision avoidance module based on artificial potential fields [14].

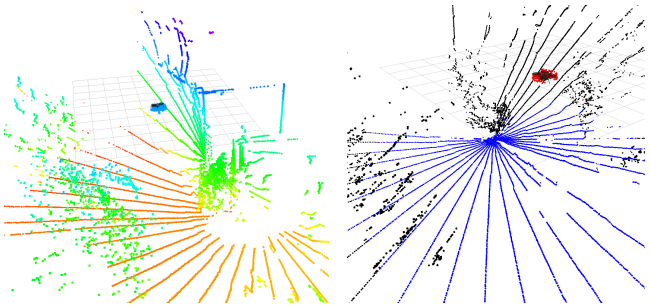

Fig. 8: Left: Acquired laser scans (points colored by height). Right: laser-based ground plane and height estimation. After filtering out measurements on the robot (red points), we estimate the dominant horizontal plane (blue points), and compute its distance to the robot as an height estimate. Obstacles other than the ground are colored black.

We have chosen this approach as a safety measure reacting directly on the available sensor information at a higher frequency than used for planning. This enables the MAV to immediately react to perceived obstacles in its vicinity and to deviations from the planned path caused by external influences such as wind. Furthermore, collision avoidance assists a human pilot to operate the MAV safely in challenging situations, e.g. flying through a narrow passageway (see Fig. 9).

As input for our algorithm, we consider the robot-centered 3D occupancy grid, the current motion state x_t, and a target waypoint w_t on a globally planned 3D path.

In the artificial potential field approach, the perceived obstacles induce repulsive forces on the flying robot. The magnitude of the repulsive force F_r of an obstacle o at a position p is calculated as $F_r^p = \text{costs}\left(argmin_o\left(\|o-p\|\right)\right)$. The MAV is directed to the target waypoint by an attractive force towards that goal. The waypoint is selected from the global path in a way that avoids the reactive approach to get stuck in a local minimum. The resulting force at a discrete position is now the weighted sum of the attractive and repulsive forces. As most of the cells in our obstacle map are free space, we do not pre-calculate the forces for every cell, but only for cells that intersect with the robot's bounding box.

In contrast to the standard potential field-based approach, we relax the assumption that the robot is an idealized particle. We account for the shape of the MAV by discretizing it into cells of the size of our 3D grid map. The center points of these cells are individual particles to the algorithm. Hence, obstacles induce repulsive forces and the target waypoint induces an attractive force on each of these cells. Thus, multiple obstacles can induce forces on different parts of the MAV. The resulting force to the MAV is now the average of the weighted sums of the individual attractive and repulsive forces.

Standard potential field approaches assume that the motion of a vehicle can be changed immediately. To overcome this limitation, we predict the MAV's future trajectory T_t given the current dynamic state x_t and the probable sequence of

Fig. 9: Reactive collision avoidance can act as a safety layer between commands given by higher-lever planning layers or a human pilot. It can help to navigate the robot (green) in challenging situations.

motion commands $u_{t:t+n}$ for a fixed discrete-time horizon n (Fig. 10). This time horizon is tightly bound by the property that multicopters can quickly stop or change their motion. To predict the trajectory, we employ a motion model of the MAV and the estimated resulting forces along the trajectory. The magnitudes of the velocity commands are calculated according to the predicted future forces.

We model the MAV's flight dynamics as a time-discrete linear dynamic system (LDS) $x_{t+1} = Ax_t + Bu_t$ that predicts the state of the MAV at time $t+1$ given the current state estimate x_t (i.e. attitude R_t, position p_t, angular ω_t and linear velocities v_t, thrust T_t) and the user command u_t. The model matrices A and B were fitted to data captured in our motion capture system by using the least squares method.

The prediction of the future trajectory for the next n time steps is then given by

$$\begin{aligned} T_t &= p_{t:t+n} = (p_t, p_t+1, \ldots, p_{t+n}), \\ p_{i+1} &= Ax_i + Bu_i + p_i \qquad i \in [t : t+n-1], \\ u_i &= C\,\vec{\boldsymbol{F}}_{\boldsymbol{p_i}}. \end{aligned}$$

The future control commands u_i are predicted by mapping the estimated forces $\vec{\boldsymbol{F}}_{\boldsymbol{p_i}}$ at a position p_i to a control command with matrix C. If a given force threshold is exceeded at any point p_i of the trajectory, we reduce the velocity v of the MAV to

$$v_{new} = \left(\frac{1}{2} + \frac{i}{2n}\right) v_{max}.$$

Fig. 10 shows our robot model and a predicted trajectory in simulation.

VII. State Estimation

To control the MAV, we need an accurate estimate of the dynamic state of the MAV at a rate equal or higher than the control frequency. A plurality of installed sensors provides us with measurements of subsets of the state variables. Furthermore, not every sensing modality is available in every situation. We fuse these measurements to a single state

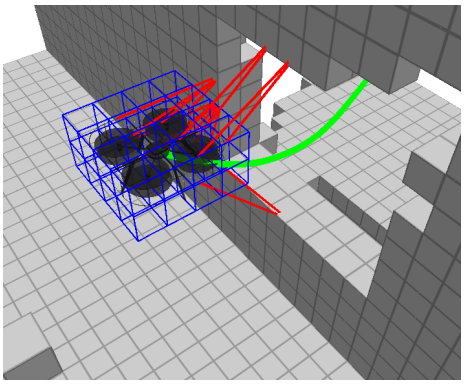

Fig. 10: We predict the influence of a motion command by rolling out the robot's trajectory (green) using a learned motion model. The current artificial repulsive forces are depicted in red.

Fig. 11: The hovering MAV (green) avoided collisions with approaching obstacles (yellow) in our experiments.

estimate using an extended Kalman Filter (EKF) based on the Bayesian Filtering Library [15].

Up to a height of 5 m, we incorporate velocity measurements from an optical flow camera (PX4FLOW [16]) at 100 Hz. Other means of velocity measurements are visual odometry using our fisheye cameras with PTAM [17] at approximately 20 Hz and GPS velocity measurements at 5 Hz coming from an u-blox LEA-6S GPS chip. This sensor is also the only source of absolute position information, if needed.

The main source for height measurements is the barometric sensor on the MikroKopter FlightControl board. In the initialization routine of the MAV, the sensor is calibrated and initialized to zero height. This sensor works under all conditions and at a high rate, but is subject to drift over time. Thus, it is mainly a good source of relative height changes. We correct these measurements with ultrasonic height measurements within the operational range of the PX4FLOW camera and with laser range measurements at 2 Hz up to a height of 30 m, as described in Sec. V-B. The attitude of the MAV is estimated using the inertial measurement unit (IMU) on the FlightControl board.

VIII. Evaluation

We evaluate the performance and reliability of our predictive collision avoidance module in simulation and on the real system.

We tested our approach on a simulated waypoint following scenario. In this scenario, the robot had to follow a path through three walls with window-like openings of different size. We measured the time the MAV needed and the forces

TABLE I: Evaluation of flight durations and repelling forces applied to the MAV.

	Time (s)	Avg. Force
Std. Potential Fields	11.9 (0.5)	0.44 (0.06)
Fixed Slow Down	12.56 (0.8)	0.43 (0.04)
Fixed Slow Down (1 s look-ahead)	14.3 (1.7)	0.28 (0.04)
Adaptive Velocities (1 s look-ahead)	12.9 (0.8)	0.3 (0.01)

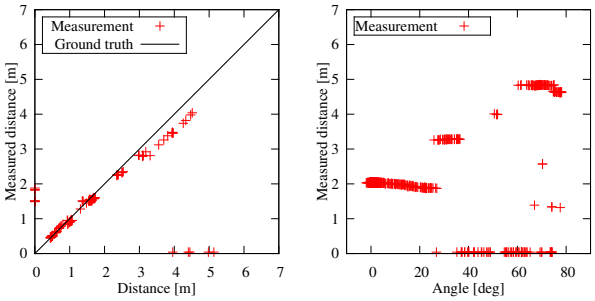

(a) Cylindrical pipe (b) Wall (varying angles)

Fig. 12: Experimental evaluation of ultrasonic range readings in different situations: varying distances to a cylindrical pipe with 20 cm diameter (left) with ground truth (black line), and distances to a wall with varying measurement angle (right).

repelling the MAV from obstacles during the flight. The forces are a measure on how close the MAV comes to the obstacles. We compared our approach with the classical potential field approach. Furthermore, we implemented a fixed slow down of the MAV. Here, the MAVs maximum speed is reduced by a fixed factor if the forces along the predicted trajectory cross a threshold at any time. In our evaluation, we tested this approach without prediction, i. e., just the forces in the current state are estimated, and with a 1 s trajectory rollout. We show the results in Tab. I.

These experiments show that our predictive collision avoidance leads to smoother trajectories, keeping the MAV further away from obstacles than the same potential field approach without trajectory prediction. No collisions occurred during these test runs. The simulated MAV was able to fly through passageways of its size plus a safety margin. We also evaluated our approach with the real robot. Our collision avoidance approach runs at approximately 100 Hz on a single core of an Intel Core 2 processor, which includes data acquisition and map building. Fig. 11 shows an experiment where the hovering MAV avoided approaching obstacles.

We evaluated our sensor setup in different situations. While the ultrasonic sensors are able to perceive relatively small objects, the measured distances might not be correct due to multiple reflections of the sent signal. Up to an angle of 25° degree from the surface normal of a wall we get correct distance measurements. Angles above that threshold induce too large distance measurements (see Fig. 12b).

Distance measurements of the laser scanner are accurate where available. Some materials, e. g., glass, are difficult to perceive. Furthermore, the sparseness of the scans makes it necessary to aggregate several scans to perceive small elongated objects. An example of the aggregated scans of a small elongated object is shown in Fig. 14.

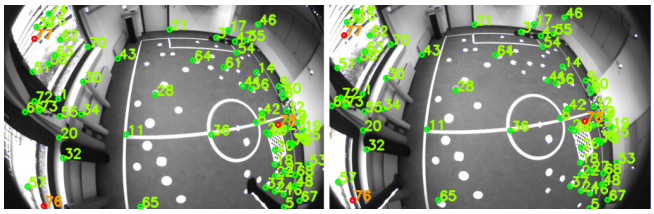

Fig. 13: Matched features of one stereo image pair.

Visual obstacle detection is based on sparse interest points, which are corners in the gradient images. The density of the triangulated obstacles depends on the texture of the obstacles, therefore we cannot detect obstacles with too less texture. Using the epipolar constraint and a high correlation coefficient as criterion for corresponding image points leads only in rare cases to wrong matches. Fig. 13 shows the matched features of one stereo pair.

IX. Conclusions

We presented an MAV equipped with a multimodal sensor setup for robust obstacle detection using fisheye stereo cameras, a lightweight 3D laser scanner, and ultrasonic sensors. These sensing modalities have different advantages and disadvantages. Laser measurements provide accurate distance measurements around the MAV, but they are sparse and fail on glass surfaces. This sensor reliably perceives larger structures like buildings, vehicles, and persons. Hence, it is well suited as our main sensor.

To detect small obstacles in the close vicinity of the MAV, we employ ultrasonic sensors. These sensors provide a coarse detection of the MAVs surrounding at a high rate. This makes ultrasonic sensors a good choice for reactive obstacle avoidance. Visual obstacles can be perceived in large distances and at a high rate, but obstacles without texture cannot be perceived reliably. These are often walls, well perceivable by our laser scanner. Hence, the plurality of sensors increases the probability to perceive all types of obstacles reliably.

Based upon this setup, we developed a fast, reactive collision avoidance layer to quickly react on new measurements of nearby obstacles. It serves as a safety measure between higher planning layers or commands given by a human pilot and the low-level control layer of the MAV. Standard potential field approaches assume that the motion of a vehicle can be changed immediately at any position in the field. To overcome this limitation, we predict the trajectory resulting from the current dynamic state and the artificial potential field into the future. This leads to safer and smoother trajectories for a multicopter.

Small obstacles might not be perceived reliably in every scan. Furthermore, occlusions by the MAV or other obstacles make it necessary to keep track of occupied space over time. In future work, we aim at improving the aggregation of obstacles into local maps over a longer time period. This needs both a good motion estimate and fast and accurate registration of new measurements to the local map. We will achieve the first objective by fusing the information from laser-based motion estimation, visual odometry, GPS, and

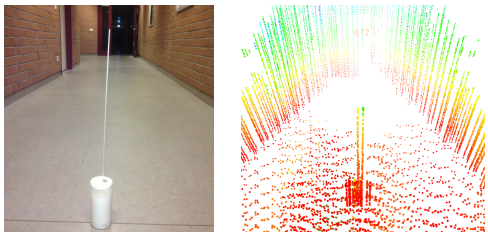

(a) (b)

Fig. 14: Perception of small obstacles with the 3D laser scanner: (a) rod with 2.5 mm diameter; (b) resulting 3D point cloud (pixel color encodes height above ground plane).

optical flow measurements into a reliable odometry estimate. Furthermore, we will develop high-level planning layers to provide globally consistent paths as input to our local obstacle avoidance.

References

[1] T. Tomić, K. Schmid, P. Lutz, A. Domel, M. Kassecker, E. Mair, I. Grixa, F. Ruess, M. Suppa, and D. Burschka, "Toward a fully autonomous UAV: Research platform for indoor and outdoor urban search and rescue," *Robotics Automation Magazine, IEEE*, vol. 19, no. 3, pp. 46–56, 2012.

[2] S. Grzonka, G. Grisetti, and W. Burgard, "Towards a navigation system for autonomous indoor flying," in *Proc. of the IEEE Int. Conf. on Robotics and Automation*, 2009.

[3] A. Bachrach, R. He, and N. Roy, "Autonomous flight in unstructured and unknown indoor environments," in *European Micro Aerial Vehicle Conf (EMAV)*, 2009, pp. 1–8.

[4] S. Shen, N. Michael, and V. Kumar, "Autonomous multi-floor indoor navigation with a computationally constrained micro aerial vehicle," in *Proc. of the IEEE Int. Conf. on Robotics and Automation*, 2011.

[5] S. Bouabdallah, P. Murrieri, and R. Siegwart, "Design and control of an indoor micro quadrotor," in *Proc. of the IEEE Int. Conf. on Robotics and Automation*, 2004.

[6] T. Puls, M. Kemper, R. Kuke, and A. Hein, "GPS-based position control and waypoint navigation system for quadrocopters," in *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, 2009.

[7] S. Grzonka, G. Grisetti, and W. Burgard, "A fully autonomous indoor quadrotor," *IEEE Trans. on Robotics*, vol. 28, no. 1, pp. 90–100, 2012.

[8] J. Shi and C. Tomasi, "Good Features to Track," in *9th IEEE Conference on Computer Vision and Pattern Recognition (CVPR'94)*, 1994, pp. 593–600.

[9] J.-Y. Bouguet, "Pyramidal Implementation of the Lucas Kanade Feature Tracker Description of the algorithm," 2000.

[10] J. Schneider and W. Förstner, "Bundle Adjustment for Omnidirectional Camera Systems with Points at Infinity Including System Calibration," *Z. f. Photogrammetrie, Fernerkundung und Geoinformation*, vol. 4, 2013.

[11] S. Abraham and W. Förstner, "Fish-Eye-Stereo Calibration and Epipolar Rectification," *ISPRS J. of Photogrammetry & Remote Sensing*, vol. 59, no. 5, pp. 278–288, 2005.

[12] S. Abraham and T. Hau, "Towards Autonomous High-Precision Calibration of Digital Cameras," in *Videometrics V, Proc. of SPIE Annual Meeting 3174*, 1997, pp. 82–93.

[13] P. H. S. Torr and A. Zisserman, "Mlesac: a new robust estimator with application to estimating image geometry," *Comput. Vis. Image Underst.*, vol. 78, no. 1, pp. 138–156, Apr. 2000.

[14] S. Ge and Y. Cui, "Dynamic motion planning for mobile robots using potential field method," *Autonomous Robots*, vol. 13, no. 3, pp. 207–222, 2002.

[15] K. Gadeyne, "BFL: Bayesian Filtering Library," http://www.orocos.org/bfl, 2001.

[16] D. Honegger, L. Meier, P. Tanskanen, and M. Pollefeys, "An open source and open hardware embedded metric optical flow cmos camera for indoor and outdoor applications," in *Proc. of the IEEE Int. Conf. on Robotics and Automation*, 2013.

[17] G. Klein and D. Murray, "Parallel tracking and mapping for small AR workspaces," in *Proc. Sixth IEEE and ACM International Symposium on Mixed and Augmented Reality (ISMAR'07)*, 2007.