

ChimneySpector: Autonomous MAV-based Indoor Chimney Inspection Employing 3D Laser Localization and Textured Surface Reconstruction

Matthias Nieuwenhuisen, Jan Quenzel, Marius Beul, David Droschel, Sebastian Houben, and Sven Behnke

Abstract—Inspection of industrial chimneys and smoke pipes induces high costs due to production downtimes and imposes risks to the health of human workers due to high temperatures and toxic gases. We aim at speeding up and automating this process with sensors mounted on multicopters. To acquire high quality sensor data, flying close to the walls of the chimney is inevitable, imposing high demands on good localization and fast and reliable control.

In this paper, we present an integrated chimney inspection robot based on a small lightweight flying platform, well-suited for maneuvering in narrow space. For navigation and obstacle avoidance, it is equipped with a multimodal sensor setup including a lightweight 3D laser scanner, stereo cameras, and a high-resolution camera for surface inspection. We tested our system in a mock-up chimney modeling several surfaces found in real chimneys, and present results from autonomous flights and the reconstruction of the chimney surface.

I. INTRODUCTION

Industrial chimneys must be inspected regularly, which causes costly production downtime—the chimney has to cool down and toxic gases have to dissipate before humans can enter—and poses dangers to human inspection personnel working in large heights. Access to higher parts of the chimney is difficult and requires cranes or scaffolding. Multicopter micro aerial vehicles (MAVs) can carry inspection sensors, such as cameras, to high altitudes. Their manual control inside chimneys close to the inspected surface is not feasible, though, especially at higher altitudes, as human pilots can hardly assess the exact distance to the surface or the MAV orientation from the ground. Furthermore, turbulences close to the walls require continuous adjustments to the control inputs, making the control even more challenging. Hence, we aim at fully autonomous inspection of chimney interiors with an MAV to make aerial platform-based inspection of chimneys feasible.

The foremost task of the autonomous system, and key functionality to build upon, is to safely navigate in the chimney in reasonable proximity to its walls. Starting from a coarse geometric model, a detailed 3D model of the chimney for localization and navigation is built by registering and aggregating measurements of a 3D laser scanner which is carried by the MAV. By fusing different sensor modalities, the MAV localizes with respect to this model.

The MAV is equipped with a high-resolution camera sensor to capture the surface of the chimney. To simplify



Fig. 1. Autonomous chimney inspection. Our lightweight chimney inspection MAV navigates autonomously assisted by real-time 6D laser localization and fast visual odometry. A double spiraling scan pattern facilitates good coverage and loop closing for camera-based surface reconstruction of the chimney walls. Our MAV is evaluated in a chimney mock-up resembling wall structures found in real industrial chimneys.

the inspection task for the human expert, we create a highly detailed model of the whole chimney. We employ reconstruction of structure from motion on the RGB images to create a detailed 3D model of the chimney surface and visualize it as unwrapped high-resolution orthoimage for the inspection by an expert. An expert can specify poses for further inspection and the MAV navigates to these poses autonomously based on the pose and velocity estimates and the 3D chimney model.

The use of autonomous MAVs will 1) diminish the risk for human inspectors, 2) reduce the costs of inspections due to shorter production downtimes, and 3) deliver high quality inspection results. We demonstrate the applicability of our system in a chimney mock-up of about 4 m in height and 3.5 m in diameter, shown in Fig. 1. Navigating and maneuvering in this very constrained space is highly challenging. Furthermore, the air flow caused by the MAV is reflected from the chimney walls and the ground and causes turbulences.

Our main contribution is an integrated lightweight MAV, including 3D laser-based 6D localization and 3D surface reconstruction, with a total mass of less than 3.5 kg. The constrained space in chimneys requires reliable navigation close to the walls—a highly demanding task given the lower computational power of the lightweight MAV compared to our prior work [1]. To the best of our knowledge, this is the first autonomous MAV-based chimney inspection system.

This work was partially funded by the European Commission in the FP7 project EuRoC (grant 608849).

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

II. RELATED WORK

Employing MAVs for inspection and surveying tasks has been an active field of research in recent years. Still, most MAVs are operated manually or via GNSS waypoint following in obstacle-free distances to objects [2], [3]. High-resolution images are taken at predefined positions for later inspection by an expert or for automated reconstruction of a model employing off-the-shelf software.

One exception is the work of Ortiz et al. [4] who developed a quadrotor MAV for autonomous vessel inspection. Similar to our approach, they employ a combination of laser localization and visual odometry for navigation, but use a system of mirrors to direct some laser beams to floor and ceiling and employ a 2D localization approach decoupled of the height measurements. In contrast, we perform full 3D SLAM and 6D localization.

A lightweight MAV system aiming at industrial boiler inspection has been presented by Burri et al. [5]. Their work focuses on agile movements in industrial environments with vision-based state estimation. Our goal is to build complete surface models of chimneys and we are, thus, interested in steady slower but accurate movements and a drift-free state estimate.

Intel demonstrated the inspection of an Airbus airplane with an MAV equipped with Intel RealSense sensors and a high-resolution camera [6]. In contrast to our work, the inspection was performed outdoors in much larger distance from the inspected surface. Furthermore, no detailed 3D model of the inspected surface employing RGB-D data was created.

In order to estimate depth of object points instantaneously for obstacle avoidance and state estimation, stereo cameras are used on MAVs, e.g., in the works of Schmid et al. [7] and Park and Kim [8]. Tripathi et al. [9] use them for reactive collision avoidance. The limited field of view (FoV) of cameras poses a problem when flying in constrained spaces like chimneys where obstacles are necessarily close to the MAV at all times.

To overcome these limitations, some MAVs are equipped with multiple (stereo) cameras. Moore et al. [10] use a ring of small cameras to achieve an omnidirectional view in the horizontal plane, but rely on optical flow for velocity control, centering, and heading stabilization only.

Grzonka et al. [11] use a 2D laser scanner to localize the MAV in environments with structures in flight altitude and to avoid obstacles. This limits obstacle avoidance to the measurement plane of the laser scanner.

Other groups combine laser scanners and visual obstacle detection [12], [13], [14]. Still, their perceptual field is limited to the apex angle of the stereo camera (facing forward), and the mostly horizontal 2D measurement plane of the scanner. They do not perceive obstacles above or below this region or behind the vehicle.

In chimney inspection, ascending flights are a main direction of movement and the MAV is operating close to surfaces such that omnidirectional obstacle perception is



Fig. 2. Sensor setup. Our MAV is a hexarotor equipped with a rotating 3D laser scanner for localization and obstacle avoidance, a stereo camera system for visual odometry estimation, and a high-resolution camera for surface reconstruction. For better illumination and shorter exposure times, the scanned wall is illuminated by bright LEDs. The 2D laser beam rotates around the yellow axis yielding laser scan lines. The laser scanner rotates around the green axis to accumulate scan lines to full 3D scans.

required. We use a continuously rotating laser scanner with a spherical field of view (FoV) that does not only allow for capturing 3D measurements without moving, but also provides omnidirectional obstacle sensing at comparably high frame rates (2 Hz in our setup).

The proposed MAV extends our own previous work [1], an MAV with a dual 3D laser scanner and three wide-angle stereo camera pairs. Another MAV with a sensor setup that allows omnidirectional obstacle perception is described by Chambers et al. [15]. In contrast to our prior work, the new MAV is significantly smaller and thus more constrained regarding payload. It has less compute power—this puts higher demands on more efficient processing—and only one stereo camera pair instead of three.

III. SYSTEM SETUP

Our chimney inspection robot is based on the new Ascending Technologies Neo hexacopter platform. With a diameter of about only 80 cm, the platform is well-suited for indoor flights. Fig. 2 shows our MAV and the used sensor setup.

The platform is equipped with a front-facing Skybotix VI-Sensor [16] used as stereo camera system for visual odometry. Each camera has a wide-angle lens with an apex angle of 122° and a resolution of 752×480 pixels. Images are captured with 20 Hz with hardware synchronization. The VI-Sensor is equipped with an IMU calibrated w.r.t. the two cameras.

On the top of the MAV, we installed a continuously rotating Hokuyo UST-20LX laser-range finder for localization and obstacle avoidance. The sensor rotates at a frequency of 1 Hz yielding a spherical 3D FoV. Due to the 270° apex angle and its mounting pose, it covers the space above the MAV with 2 Hz—chimney inspection starts at the ground and thus unknown obstacles are more likely to be above the MAV—and the space below with 1 Hz.

For surface reconstruction and inspection, we mounted an Intel RealSense SR300 RGB-D camera [17] at the rear-end of the MAV. The camera is mounted on a 32 cm long cantilever to take the sensor closer to the surface while mitigating the

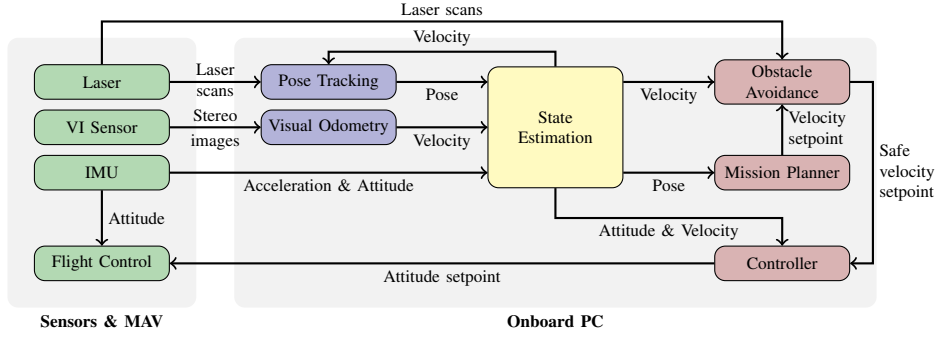


Fig. 3. Schematic overview of system components. Measurements from laser scanner and stereo cameras are processed to pose and velocity estimates. We filter these estimates with IMU measurements to get the MAV state with low-latency at 100 Hz. A mission planner generates and executes chimney coverage tours or plans paths to all inspection poses for targeted inspection.

effects of turbulences when flying close to walls. For surface reconstruction, we use the rolling shutter RGB camera of the sensor running at 30 Hz. To avoid strong motion blur during dynamic MAV motion—the necessary exposure time for the RGB camera to obtain well-illuminated images is quite high—we added two stripes of bright LEDs on the sensor. This allows us to set the exposure time to 8 ms, resulting in still dark images that are sufficiently sharp to be used for reconstruction.

All navigation-relevant sensor data processing, localization, and planning is performed onboard the MAV. Thus, the MAV can operate fully autonomously without depending on an unreliable WiFi connection to a ground control station. To facilitate this, the MAV is equipped with a small and lightweight Intel NUC PC with Intel Core i7-5557U dual core CPU running at 3.1 GHz and 16 GB of RAM. For surface reconstruction, we require high-resolution images. Consequently, the full HD SR300 sensor generates 200 MB of data per second with our settings. To store this data stream, we equipped the onboard PC with a fast Samsung Pro 950 SSD.

The 3D laser scanner and the Skybotix VI-Sensor are connected via Ethernet, and the laser rotator via a serial-to-USB interface. The Intel RealSense SR 300 RGB-D camera is connected via USB 3.0 to the onboard PC.

Fig. 3 gives a schematic overview of the flight-relevant system components. The overall weight of the system with all sensors and batteries is about 3.4 kg.

IV. LOCALIZATION AND STATE ESTIMATION

In order to navigate in a chimney, robust localization and state estimation not relying on GNSS availability are crucial. Our multimodal localization and state estimation pipeline exploits the specific characteristics of all sensors in terms of, e.g., accuracy and speed.

Visual Odometry

Our visual odometry estimation is based on LIBVISO2 [18], a fast feature-based visual odometry library for monocular and stereo cameras. The stereo odometry approach is very general and does not require a specific motion model. The only prerequisite is that the input images

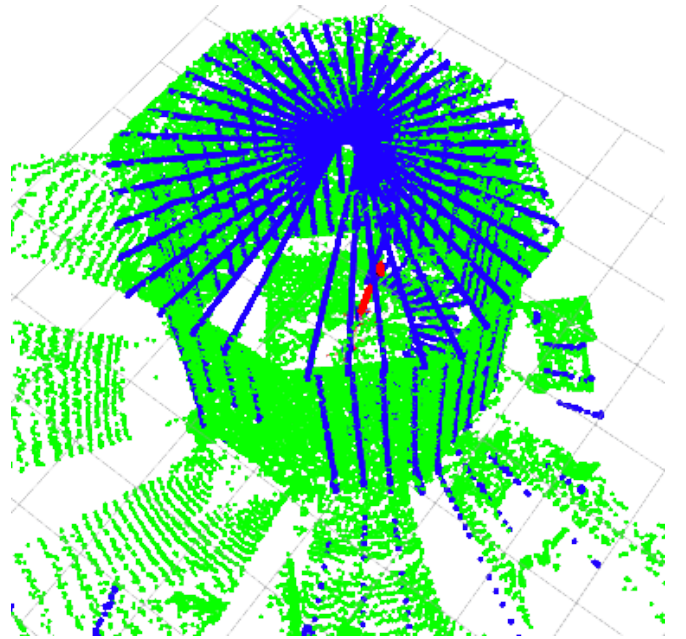


Fig. 4. Laser-based localization. To track the MAV pose (red arrow) in an allocentric frame, we aggregate individual laser scans over 500 ms (blue) and match these to an allocentric map of the chimney (green).

have to be rectified and the calibration parameters have to be known. Similar to other feature-based methods, LIBVISO2 extracts and matches features over subsequent stereo frames and estimates the egomotion by minimizing the reprojection error. To be robust to outliers, RANSAC is used for initialization of the minimization step.

The interior of a chimney poses a challenging environment for visual odometry algorithms—the tracked features are always close to the camera due to the restricted space and the repetitive structures are self-similar. Thus, visual odometry is prone to heavy drift and can easily lose track. We use only the position derivative of the visual odometry and integrate it as noisy velocity estimate into our state estimation filter at the stereo camera frequency of 20 Hz.

To localize the MAV in an allocentric chimney frame, we track its pose by registering local multiresolution maps to a global map employing multiresolution surfel registration (MRSR) [19]. This yields a 6D pose estimate in the map frame at 2 Hz (Fig. 4). We build the map of the chimney ad-hoc from the takeoff position before a mission. For larger chimneys, our approach is able to perform SLAM during a first simple exploration flight, e.g., flying straight up and down in the chimney center. The allocentric map is represented by surfels with a uniform size.

Since the laser scanner acquires complete 3D scans with a relatively low frame rate, we incorporate the filtered egomotion estimate from visual odometry and measurements from the Neo’s 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 filter result and alignment of the current local multiresolution map towards the allocentric map of the environment.

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 only track one pose hypothesis. We assume that the initial pose of the MAV is roughly known by starting from a predefined pose in the center of the chimney or, if this is not possible, setting the pose in our control GUI. The approximately known pose is then quickly refined by scan registration. Here, small structures, e.g., a ladder commonly attached to chimney walls, are sufficient to align with the map.

Fig. 4 shows the registration of a 3D scan to the map and an estimated 6D pose. The resulting robot pose estimate from the allocentric localization is used as a measurement update in a lower-level state estimation filter.

For obstacle avoidance, we create a local obstacle map consisting of the most recent individual laser measurements that form a full laser rotation, i.e., 1 s of aggregated laser scans, depicted in Fig. 5. The obstacle map is kept in an egocentric frame by incorporating the MAV egomotion estimate.

State Estimation Filter

We use two filters for state estimation: A low-level filter fuses measurements from accelerometers, gyros, and compass to one 6D attitude and acceleration estimate in the AscTec Trinity flight control unit. The second higher-level filter fuses linear acceleration, velocity, and position information to a state estimate that includes 3D position and velocities. The higher-level filter is based on the Pixhawk Autopilot [20] position estimator adapted to use visual odometry velocity estimates and laser pose tracking. IMU measurements are incorporated at 100 Hz and a pose/velocity estimate is published at the same rate.




Fig. 5. Laser obstacle map. To avoid collisions, we maintain an obstacle map containing the most recent 3D laser scan and move it with our motion estimate (black arrow). The blue arrow depicts the laser localization pose.

In our position filter, we estimate the state

$$x = \begin{pmatrix} p_x & p_y & p_z \\ v_x & v_y & v_z \\ a_x & a_y & a_z \end{pmatrix},$$

consisting of position p , velocity v , and acceleration a . The prediction step, based on the current bias-corrected acceleration measurement $a_{k,sens}$, is

$$\begin{aligned} p_k &= p_{k-1} + v_{k-1} \cdot dt + \frac{1}{2} a_{k,sens} \cdot dt^2, \\ v_k &= v_{k-1} + a_{k,sens} \cdot dt, \\ a_k &= a_{k,sens}. \end{aligned}$$

If sensor measurements are available, i.e., the measurement is not timed out, the state is corrected accordingly. For velocity measurements $v_{k,sens}$, coming from, e.g., visual odometry, the state correction is:

$$\begin{aligned} v_k &= v_{k-1} + (v_{k,sens} - v_{k-1}) \cdot w \cdot dt, \\ a_k &= a_{k-1} + (v_{k,sens} - v_{k-1}) \cdot w^2 \cdot dt^2. \end{aligned}$$

Here, w is a weighting factor that indicates the reliability of the inputs.

Position measurements $p_{k,sens}$, coming from, e.g., the laser scanner, are incorporated as

$$\begin{aligned} p_k &= p_{k-1} + (p_{k,sens} - p_{k-1}) \cdot w \cdot dt, \\ v_k &= v_{k-1} + (p_{k,sens} - p_{k-1}) \cdot w^2 \cdot dt, \\ a_k &= a_{k-1} + (p_{k,sens} - p_{k-1}) \cdot w^2 \cdot dt^2. \end{aligned}$$

If no new sensor measurements are received, we gradually degrade the velocity estimate in the correction step until the filter stops.

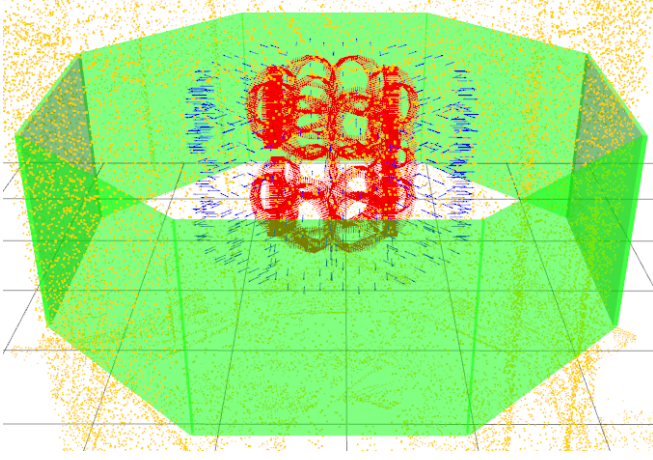


Fig. 6. Coverage tour in the chimney. We plan sensor coverage tours in a coarse geometric chimney model (green) aligned with the allocentric laser map (yellow). The sensor is moved in circles to reach a good coverage of the surface and to facilitate loop closing (blue arrows). The corresponding MAV path to achieve the desired sensor poses is depicted with red arrows.

V. PLANNING AND NAVIGATION

Capturing the surface of the chimney with our RGB-D camera requires a steady flight path with a fixed distance between sensor and walls. Furthermore, the images need sufficient overlap in every direction to build a consistent model for the whole flight. These demands are hard to fulfill in manual operation, especially given the turbulent air movement close to the walls pushing the MAV away and requiring constant control actions. Thus, we operate the MAV fully autonomously except for start and landing. First, we plan an inspection path with 50 % image overlap starting 1 m above the ground and moving upwards in a spiraling motion. Along this primary spiral, the MAV flies smaller vertical circles to ensure a good image overlap and loop closures. This pattern has been found to be advantageous over a simple spiraling motion in preliminary flight tests. The inspection mission is planned, given a simple geometric model of the chimney, the sensor characteristics (apex angles, best scanning distance), and the part of the chimney to cover. Fig. 6 shows an example inspection path in our octagonal mock-up chimney.

After a first complete inspection, the user can specify poses for a targeted second inspection, e.g., to take close-up images of potential defects in the chimney. The MAV processes a set of inspection poses and determines an optimal processing order to achieve a short inspection flight employing a traveling salesman problem solver (Fig. 11). Selection of poses for targeted inspection is assisted by a graphical tool on the ground station that shows the taken images to an operator and determines corresponding MAV poses for selected images. Thus, the operator can select images where a possible defect is visible and mark these poses for the second inspection.

To safely navigate in the vicinity of obstacles, e.g., probes, ladders, or open hatches in the chimney, we employ reactive




Fig. 7. Scaling factors for reactive obstacle avoidance. If the MAV is too close to an obstacle, it is actively pushed away by a force scaled with s_{push} . Entering a safety zone around obstacles is prevented by reducing velocities towards the obstacle with factor s_{slow} .

local obstacle avoidance. For this, we extended our previous work on reactive obstacle avoidance [21]. Based on a laser map of the vicinity of the MAV (see Fig. 5), we reduce the MAV velocities towards close obstacles and actively push the MAV back from obstacles if the distance to an obstacle falls below a safety distance. The resulting velocity command v_c given a target velocity v_t and the artificial force F induced by an obstacle is calculated as

$$v_o = (v_t \cdot \hat{F}) \hat{F}, \quad (1)$$

$$v_c = v_t - s_{\text{slow}} v_o + s_{\text{push}} \hat{F}, \quad (2)$$

where v_o is the part of the target velocity v_t towards the obstacle. Negative parts of v_o , i.e., velocities maneuvering the MAV away from the obstacle, are set to zero. The scalar factors s_{slow} and s_{push} denote the strength of slowing down the MAV and actively pushing it away from obstacles (see Fig. 7). Their value is a linear interpolation between free-space distance and the safety distance, and an interpolation between safety distance and critical distance, respectively. This extends our previous approach on reactive obstacle avoidance to allow for less conservative safety distances, as required by our RGB-D sensor, while still maintaining safe navigation.

Velocity setpoints for the MAV are generated by means of a PID-controller in the integrated mission planning and navigation node when executing coverage or inspection missions. These setpoints are input to our obstacle avoidance module. After processing these setpoints, the resulting safe velocity setpoints are fed to the velocity controller to generate attitude and thrust commands. Our velocity controller is based on the linear MPC controller by Kamel et al. [22], modified to work with velocity-only setpoints. The resulting attitude-thrust setpoint is controlled by the low-level AscTec Trinity Autopilot of the MAV.

VI. SURFACE RECONSTRUCTION

The objective of our inspection flight is a high-resolution surface reconstruction of the chimney walls. The reconstruction is based on the data acquired with the SR300 sensor at the rear end of the MAV. For the inspection, it is necessary that also smaller details, e.g., cracks in the wall, are correctly modeled in the reconstruction shown to inspection experts. Thus, we use the full 1080p resolution of the RGB camera of the sensor. Our results are a colored high-resolution mesh




Fig. 8. Schematic overview of surface reconstruction. We derive an accurate camera trajectory, estimate the 3D structure of the surface, and fuse the data to colored meshes.

and an unwrapped 2D orthoimage shown to an expert for visual inspection. Fig. 8 shows an overview of our surface reconstruction pipeline.

Whereas the MAV pose estimated in real-time during flight is accurate enough for control and navigation, the surface reconstruction demands poses with much higher precision. Thus, in order to employ a mapping with known poses approach, the accurate MAV trajectory has to be estimated in an offline processing step making use of RGB image data from the SR300 sensor. To estimate the trajectory, we employ the structure-from-motion software COLMAP [23]. Instead of using each image for sparse reconstruction, we decided to use only two images per second to keep the calculation within a reasonable timespan. We then export the sparse pointcloud and the trajectory into MVE [24] and perform dense reconstruction per multi-view stereo at full image resolution. The dense reconstruction contains holes especially in dark regions with low intensity gradients. We employ the Domain Transform [25] to interpolate depth values in these regions with the RGB images as a guide. The filled depth maps are then backprojected by MVE to create a dense point cloud. The mesh is generated using the Floating Scale Surface Reconstruction [26]. After sub-sampling the mesh multiple times, we use Color Map Optimization [27] to refine the image poses with respect to the mesh such that the intensity difference between all images observing a mesh-vertex becomes minimal. Instead of the restricted image set used for sparse reconstruction, we use now all captured images. The image poses are initialized by interpolation from the reconstruction poses. The final mesh color is obtained as a weighted average of the mesh-vertex observations. A section of the generated high-resolution mesh is shown in Fig. 9. Finally, the reconstructed mesh is unwrapped to yield an easily inspectable 2D visualization, depicted in Fig. 10. This is done by finding a specified number of planes within the 3D model.

VII. RESULTS

To test our integrated system, we inspected an octagonal chimney mock-up, shown in Fig. 1. The mock-up was designed by a chimney inspection service contractor to facilitate the transferability to real inspection applications. The chimney is constructed from wooden panels with styrofoam structures on the inner sides resembling the stonework and concrete patterns that can be found in many industrial chimneys. Single structure elements are of size 1×0.5 m and each of the stone walls—except of the wall containing




Fig. 9. Comparison of reconstructed surface with a photograph. Close up of the rightmost chimney wall in our reconstruction depicted in Fig. 10 (left) and image of the surface taken with a compact camera (right). The defects in the wall are clearly visible in the reconstruction (circled yellow).

an entry to the mock-up—is plastered with a single type of elements resulting in repetitive patterns.

In addition, one panel carries a rusty iron surface as found in chimneys with a metal alloy on the inner side. Bricks were carved out of the styrofoam to represent defects. The chimney mock-up is 4.4 m high and the inradius is 1.8 m. With a required sensor distance of 0.8 m to 1.0 m to the surface, the remaining safe navigation space has a diameter of approximately 1.2 m.

We started with an initial coverage flight to acquire RGB data of the chimney surface. The flight to acquire the data used for the reconstruction shown in Fig. 10 was seven minutes. Only start and landing phase were controlled manually. For surface coverage, the MAV followed a horizontally and vertically spiraling pattern with the RGB-D sensor directed to the nearest surface in order to enable loop closings in the later surface reconstruction. All defects were covered in that time.

After the coverage flight, we downsample the recorded video stream from the RGB-D sensor to speed-up the process of transferring it to a ground control station. Here, an operator identifies poses for a more detailed inspection in the video stream and marks them. In the test case, the operator could identify all ten defects in the images. The MAV poses corresponding to the images showing defects were stored for a second flight to reinspect those defects. To exemplify the targeted inspection of previously identified defects, the MAV then autonomously planned an inspection mission and followed a path adopting all stored poses in a



Fig. 10. 2D visualization of the chimney walls. Reconstructed surface of seven walls from the mock-up chimney shown in Fig. 1. Despite some warping effects caused by the unwrapping of the 3D structure, all details are accurately reconstructed.

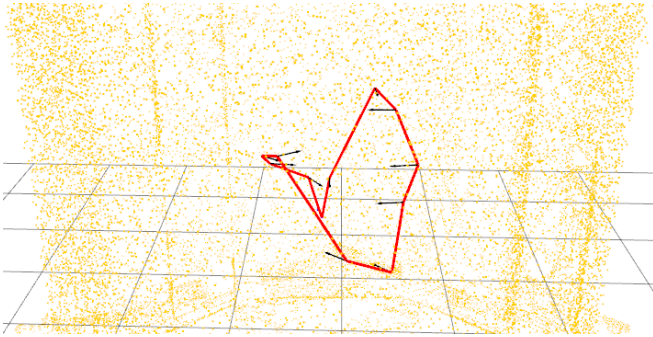


Fig. 11. Targeted inspection. After an operator selected targets to reinspect, the observation poses (black arrows) are sent to the MAV where a mission planner finds an optimal processing order (red path). The MAV navigates to these poses autonomously and hovers there for several seconds to acquire more detailed data of the surface. The MAV map is depicted by yellow dots.

useful order (Fig. 11) and held its position for several seconds to demonstrate reaching the desired position. All poses were successfully reached and the defects were clearly identified in the data captured during the second flight. After the targeted inspection, the acquired data was transferred to the ground control station for further offline processing.

A video showing an initial coverage flight in the mock-up chimney and the following targeted inspection can be found on our website¹. On average, we acquired depth information for 95.5 % of the pixels, thus, the sensor was almost always positioned in an optimal distance to the surface.

In a post-processing step, we reconstruct the surface of the chimney walls. Fig. 10 shows the resulting unwrapped 2D visualization. We could successfully reconstruct all of the eight segments. All segments covered with artificial stonework were reconstructed without major artifacts. Overall, the representation encompasses all interesting details, in particular, all ten defects carved into the styrofoam wall. The reconstruction of the rusty metal segment was sufficient for visual inspection. Depth could not be estimated accurately for some parts of the metal wall as the images were too dark to yield enough features—a result of the relatively low exposure time required to avoid motion blur. This issue can

be addressed by either an other high-resolution camera or additional onboard illumination.

The computational load during the coverage flight posed a particular challenge due to the laser and visual odometry-based state estimation on the two available physical CPU cores. In combination with recording of the data streams from both the SR300 and stereo cameras desired for optimal post-processing, navigation and MAV control showed an unstable behavior during initial tests. Hence, we omitted to record all non-crucial data—including the stereo cameras—and estimated the MAV trajectory for surface reconstruction solely from the the monocular vision data from the SR300 sensor. Furthermore, we capture frames at a reduced frequency of 30 Hz instead of the desirable maximum frequency of 60 Hz. Thus, post-processing is more challenging due to the necessary compensation for larger inter-frame motions. Still 200 MB s^{-1} were recorded from the SR300 sensor during the coverage flight. With these means of reducing the system load and by tuning the onboard PC to reach a steady data flow without peaks, we accomplished stable and safe system performance.

The 3D laser scanner served well for low-frequency allocentric localization in a previously recorded map in combination with visual odometry from the VI stereo sensor to compensate for the MAV motion between scans. Furthermore, this combination yielded a high-frequency positioning and state estimation and allowed for stable MAV control when filtered with IMU readings. The stability of this approach allowed us to rely on a pre-chosen initial pose in order to resolve the orientation ambiguity by the symmetric chimney cross section.

We evaluated the reactive obstacle avoidance system quantitatively as part of the European Robotics Challenge, combined with stereo camera obstacle perception instead of the laser obstacle map. The MAV was steered velocity controlled through free-space and towards obstacles of different shape and size. The maximum linear velocity was 1 m s^{-1} . Our MAV was able to safely stop in front of every obstacle. In 40 % of the evaluated cases in the optimal distance range, in the other cases not more than 0.4 m off, yet still in a safe range. This error is the combined error of visual perception

¹http://www.ais.uni-bonn.de/videos/ICUAS_2017_ChimneySpector

and reactive obstacle avoidance.

VIII. CONCLUSIONS

In this work, we have presented an integrated chimney inspection system based on a lightweight MAV platform. The small form factor severely restricts payload. Thus, choosing lightweight sensors and a small onboard PC with limited compute power is essential. To this end, we reduced the onboard sensor setup from our previous work to a smaller subset and still maintain good real-time state estimation and obstacle perception. In addition, we added a new sensor for surface reconstruction. Integrating the whole sensor, state estimation, and control setup on a lightweight MAV posed challenges, but could be successfully achieved. Important lessons learnt are, that special care has to be taken on required data for online and offline processing and a sensible assessment of the required accuracies and real-time performance for every part of the system has to be performed. Overall, the MAV can cover chimney surfaces autonomously and a high-quality surface reconstruction is possible in a post-processing step. Defects identified by experts can be reinspected at more detail in a second mission.

ACKNOWLEDGMENTS

The authors wish to thank the Autonomous Systems Lab from ETH Zürich and Ascending Technologies for their technical and organizational support, especially for providing the flying platform. We wish to thank CRN Management GmbH for building the mock-up chimney, for providing expertise for the domain of chimney inspection, and for their support during the evaluation flights.

REFERENCES

- [1] M. Nieuwenhuisen, D. Droschel, M. Beul, and S. Behnke, "Autonomous navigation for micro aerial vehicles in complex GNSS-denied environments," *Journal of Intelligent and Robotic Systems*, vol. 84, no. 1, pp. 199–216, 2016.
- [2] F. Nex and F. Remondino, "UAV for 3D mapping applications: a review," *Applied Geomatics*, vol. 6, no. 1, pp. 1–15, 2014.
- [3] B. Chan, H. Guan, J. Jo, and M. Blumenstein, "Towards UAV-based bridge inspection systems: a review and an application perspective," *Structural Monitoring and Maintenance*, vol. 2, no. 3, pp. 283–300, 2015.
- [4] A. Ortiz, F. Bonnin-Pascual, and E. Garcia-Fidalgo, "Vessel inspection: A micro-aerial vehicle-based approach," *Journal of Intelligent and Robotic Systems*, vol. 76, no. 1, pp. 151–167, 2014.
- [5] M. Burri, J. Nikolic, C. Hürzeler, G. Caprari, and R. Siegwart, "Aerial service robots for visual inspection of thermal power plant boiler systems," in *Proc. of Int. Conf. on Applied Robotics for the Power Industry (CARPI)*, 2012.
- [6] Intel Corp., "Intel and Airbus demo drone inspection of passenger airliners," <https://newsroom.intel.com/chip-shots/intel-airbus-demo-drone-inspection-of-passenger-airliners/>, 2016.
- [7] 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.
- [8] J. Park and Y. Kim, "3d shape mapping of obstacle using stereo vision sensor on quadrotor uav," in *AIAA Guidance, Navigation, and Control Conf.*, 2014.
- [9] A. Tripathi, R. G. Raja, and R. Padhi, "Reactive collision avoidance of UAVs with stereovision camera sensors using UKF," in *Advances in Control and Optimization of Dynamical Systems*, 2014, pp. 1119–1125.
- [10] R. Moore, K. Dantu, G. Barrows, and R. Nagpal, "Autonomous MAV guidance with a lightweight omnidirectional vision sensor," in *Int. Conf. on Robotics and Automation (ICRA)*, 2014.
- [11] S. Grzonka, G. Grisetti, and W. Burgard, "A fully autonomous indoor quadrotor," *IEEE Trans. on Robotics*, vol. 28, no. 1, pp. 90–100, 2012.
- [12] T. Tomić, K. Schmid, P. Lutz, A. Domel, M. Kassecker, E. Mair, I. Gríxa, 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.
- [13] S. Huh, D. Shim, and J. Kim, "Integrated navigation system using camera and gimbaled laser scanner for indoor and outdoor autonomous flight of UAVs," in *Int. Conf. on Intelligent Robots and Systems (IROS)*, 2013, pp. 3158–3163.
- [14] B. Jutzi, M. Weinmann, and J. Meidow, "Weighted data fusion for UAV-borne 3D mapping with camera and line laser scanner," *Int. J. of Image and Data Fusion*, 2014.
- [15] A. Chambers, S. Achar, S. Nuske, J. Rehder, B. Kitt, L. Chamberlain, J. Haines, S. Scherer, and S. Singh, "Perception for a river mapping robot," in *Int. Conf. on Intelligent Robots and Systems (IROS)*, 2011.
- [16] J. Nikolic, J. Rehder, M. Burri, P. Gohl, S. Leutenegger, P. T. Furgale, and R. Siegwart, "A synchronized visual-inertial sensor system with FPGA pre-processing for accurate real-time SLAM," in *Int. Conf. on Robotics and Automation (ICRA)*, 2014.
- [17] Intel Corp., "Intel RealSense camera SR300 embedded coded light 3D imaging system with full high definition color camera—product datasheet," 2016.
- [18] A. Geiger, J. Ziegler, and C. Stiller, "StereoScan: Dense 3D reconstruction in real-time," in *IEEE Intelligent Vehicles Symposium*, 2011.
- [19] D. Droschel, J. Stückler, and S. Behnke, "Local multi-resolution representation for 6D motion estimation and mapping with a continuously rotating 3D laser scanner," in *Int. Conf. on Robotics and Automation (ICRA)*, 2014.
- [20] L. Meier, P. Tanskanen, L. Heng, G. Lee, F. Fraundorfer, and M. Pollefeys, "PIXHAWK: A micro aerial vehicle design for autonomous flight using onboard computer vision," *Autonomous Robots*, vol. 33, no. 1–2, pp. 21–39, 2012.
- [21] M. Nieuwenhuisen, M. Schadler, and S. Behnke, "Predictive potential field-based collision avoidance for multicopters," in *Int. Arch. Photogramm. Remote Sens. Spatial Inf. Sci. (ISPRS)*, vol. XL-1/W2, 2013, pp. 293–298.
- [22] M. Kamel, M. Burri, and R. Siegwart, "Linear vs Nonlinear MPC for Trajectory Tracking Applied to Rotary Wing Micro Aerial Vehicles," *ArXiv e-prints*, Nov. 2016.
- [23] J. L. Schönberger and J.-M. Frahm, "Structure-from-motion revisited," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [24] S. Fuhrmann, F. Langguth, and M. Goesele, "Mve—a multi-view reconstruction environment," in *EUROGRAPHICS Workshops on Graphics and Cultural Heritage*, 2014.
- [25] E. S. Gastal and M. M. Oliveira, "Domain transform for edge-aware image and video processing," *ACM Transactions on Graphics (ToG)*, vol. 30, no. 4, p. 69, 2011.
- [26] S. Fuhrmann and M. Goesele, "Floating scale surface reconstruction," in *ACM Trans. on Graphics (Proc. of ACM SIGGRAPH 2014)*, 2014.
- [27] Q.-Y. Zhou and V. Koltun, "Color map optimization for 3D reconstruction with consumer depth cameras," *ACM Trans. Graph.*, vol. 33, no. 4, pp. 155:1–155:10, 2014.