

Micro Aerial Projector - Stabilizing Projected Images Of An Airborne Robotics Projection Platform

Werner Alexander Isop¹, Jesus Pestana¹, Gabriele Ermacora², Friedrich Fraundorfer¹, Dieter Schmalstieg¹

Abstract—A mobile flying projector is hard to build due to the limited size and payload capability of a micro aerial vehicle. Few flying projector designs have been studied in recent research. However, to date, no practical solution has been presented. We propose a versatile laser projection system enabling in-flight projection with feedforward correction for stabilization of projected images. We present a quantitative evaluation of the accuracy of the projection stabilization in two autonomous flight experiments. While this approach is our first step towards a flying projector, we foresee interesting applications, such as providing on-site instructions in various human machine interaction scenarios.

I. INTRODUCTION

Major improvements in terms of how humans interact with machines and digital information are still ongoing. In the last years, interfaces based on direct-touch or devices with gesture recognition have come to maturity. Mobile portable devices, like smartphones, and wearables, like head mounted displays, are becoming widespread. However, they require either visual or physical attention and constrain the user.

Spatial augmented reality [1] tries to evade those constraints, but is strongly dependent on projection devices with significant weight, which therefore have to be considered as stationary. As a consequence, it is difficult to cover wide projection areas.

We propose to address these limitations by combining augmented reality and mobile robotics into a new form of human-machine interaction. Specifically, we introduce a small semi-autonomous micro aerial vehicle (MAV) with an onboard lightweight laser-projection system and visual sensors, called the Micro Aerial Projector (MAP) (Fig. 2). We think of it as a robotic companion, which follows the user and is able to project supportive information in the 3D environment. Fig. 1 shows an example where the MAP assists a student solving mathematical problems by projecting results into the environment.

The design of the MAP requires mastering several challenges. First, safety considerations require that the MAP has to be small sized and as lightweight as possible. It should provide enough payload for all required input, output and computational units. Furthermore, it has to offer sufficient flight dynamics and flight times for being able to follow the user for an adequate amount of time. To meet these

¹Werner Alexander Isop, Jesus Pestana, Friedrich Fraundorfer and Dieter Schmalstieg are with the Department of Computer Science, Graz University of Technology, 8010 Graz, Austria. (email:{isop, pestana, fraundorfer, schmalstieg}@icg.tugraz.at)

²Gabriele Ermacora is with the Department of Computer Science, Politecnico di Torino, 10129 Turin, Italy. (email:gabriele.ermacora@polito.it)

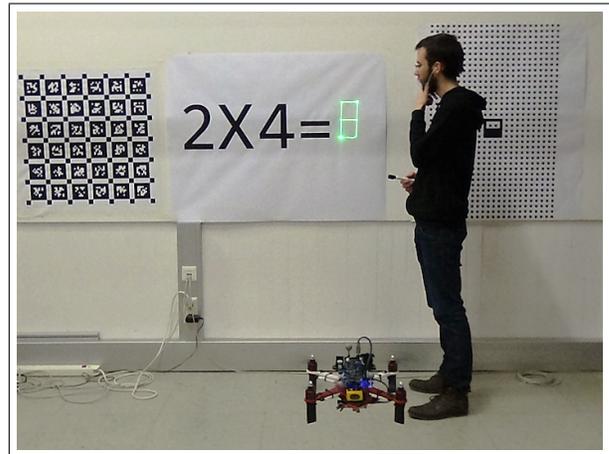


Fig. 1: The MAP supporting a student by recognizing an equation and projecting the result during hover flight.

requirements, we introduce a novel laser projection system built from scratch. It is small, lightweight and can project a set of basic symbols for a wide variety of human machine interaction (HMI) scenarios.

Second, the quality of airborne projection is influenced by the stability of the MAV. Changes to position and orientation of the drone, while hovering or during dynamic flight, as well as vibrations emerging from the rotors during flight, must be considered. Thus, a stabilization method is necessary to compensate at least for significant movements of the drone. To this aim, we propose a simple but robust feedforward correction approach, which is able to tackle image fluctuations by deflecting the projector's laser beam. The feedforward correction algorithm is based on pose estimates of an Optitrack motion tracking system. We evaluate quality of the projection stabilized by directly using the pose estimates from the tracking system and compare it to utilizing sensor fusion with the IMU, which is implemented in the inertial state estimation of the onboard flight management controller.

The contributions of this work are the following. As part of the proposed scenarios, this paper represents our first step towards combining the fields of mobile robotics, spatial augmented reality and HMI, focusing on the MAP as a small sized and potent flying projection platform. We introduce a novel lightweight projection system built from scratch, complemented by a projector calibration model. Our system is able to project a steady pattern from a moving drone. We quantitatively evaluate the accuracy of the steady pattern projection during autonomous hovering and dynamic flight.

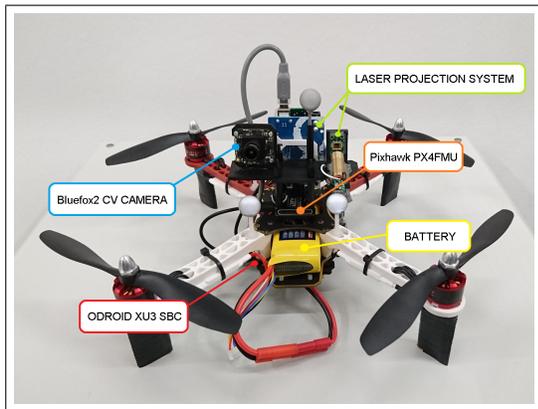


Fig. 2: MAV platform of our experimental setup including the main components. The laser projector together with the camera is mounted on top. Below, the flight management controller is located. The battery is mounted in the middle to balance weight distribution. The onboard computer is located at the bottom.

Furthermore, we examine the overall system capabilities, discussing the difficulties and limitations when putting the MAP into practice.

II. RELATED WORK

The MAP is a conceptual extension of a mobile projector. Since the early 2000s, spatial augmented reality has been a popular topic for mobile interaction. However, most approaches for room-sized environments either cover all relevant surfaces with arrays of projectors or rely on steerable projection. This idea was pioneered with the everywhere display projector [2] and later extended with real-time reconstruction from commodity depth sensors in the Beamatron project [3]. Hörtnner et al. [4] further introduced a spatial display paradigm without using a projector, but controllable moving visible objects in 3D physical space, the so called Spaxels.

Other work considers mobile projectors, which are hand-held or worn on the users head or body. They can be combined with depth sensors [5] or inertial sensors [6] to react to the user's movement and, potentially, a changing environment. However, visual light projectors with sufficient brightness for spatial interaction usually require a stationary power source. Truly mobile, battery-powered projectors can only operate at very short distances.

Laser pointers concentrate the emitted energy in a single spot and, consequently, can achieve a significantly better contrast than conventional projectors with the same power budget. A steerable laser pointer can be used to point to a particular task location. A shoulder-mounted implementation of a steerable laser pointer has been used in a tele-assistance scenario [7]. However, the same authors later proof that a single point is not sufficient for conveying complex instructions [8].

Instead of a steerable single laser point, a scanning laser can be used. Maeda et al. [9] describe a head-mounted

projective display with a scanning laser mounted co-axial to the observer's eye. Schwerdtfeger et al. [10] explore head-mounted and stationary scanning lasers for spatial augmentation. In both cases, the workspace is a major limitation.

Because of the obvious technical difficulties, there has been little work on flying projectors. Scheible et al. [11] demonstrate outdoor flying projection with a commercial visual light projector weighting 200g, mounted on a large octocopter platform with payloads of up to 3.5kg. These authors suggest a human-computer interaction scenario, but do not consider the problem of projection stabilization.

The closest work to ours in terms of combining airborne projection was presented by Hosomizo et. al [12]. They suggest a flying projection platform and approach the problem of image stabilization by combining dead reckoning and computer vision. However, they, again, use a commercial, heavy projection system and off-board computation of the image stabilization. They do not report exact measurements of the uncompensated position of the projection. Furthermore they do not provide information about the distance to the projection surface during stabilization. Moreover results for image stabilization are not evaluated in flight and only shown while the MAV is suspended from wires to overcome weight constraints.

In contrast, we introduce a small sized MAV platform, which has all essential components - a single board computer (SBC), a vision camera and a projector - included onboard. We deploy a custom lightweight steerable laser projector, which is calibrated, and also evaluate performance of projection stabilization during flight. We use inter-sensor calibration of the laser projector with respect to the flight management controller's IMU and the motion tracking system. Additionally, we also propose to describe the intrinsics of the projector with a model which is similar to a camera pinhole model. This approach not only improves accuracy of image stabilization, but also makes it compatible to common computer vision algorithms.

III. SYSTEM DESCRIPTION

This section describes our system setup including the MAP and discusses its characteristics and limitations. We use an Optitrack motion tracking system based on eight tracking cameras. The Optitrack server is connected to the groundstation via Ethernet. The groundstation is connected to the MAP through WiFi. The laser projection system and the PX4 autopilot are mounted on the MAP and interfaced to the onboard computer via serial links. Fig. 3 shows an overview of the setup.

In Fig. 2, we show the ready-to-fly MAP. The battery is mounted in the middle for balanced weight distribution. Below the frame, the on-board computer is positioned. Above the battery, the low-level flight management controller is located. The projection system and the camera are mounted on top, facing in forward flight direction. The camera is inclined to keep the rotors out of the field of view (FOV).

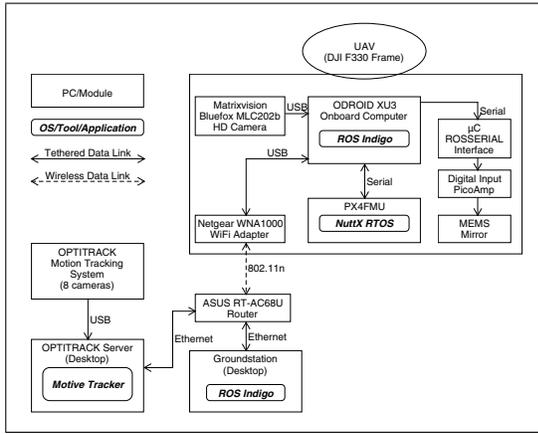


Fig. 3: The main components of our setup include the MAP itself, the on-board computer and the custom laser projector. We use an Optitrack motion tracking system combined with a groundstation to control the MAV.

A. Micro-Aerial Vehicle

The MAV, which is of 33cm frame diameter and 1200g weight, uses a semi-customized design with engines, rotors and frame taken from a DJI F330 platform. It is powered by a single 2700mAh battery with 14.8V. The flight time is acceptable with about 11 minutes, while running all relevant components and tasks for stabilization of the projection. As an onboard SBC we use an ODROID XU3 with a Samsung Exynos 5422 processor. We added a Pixhawk PX4 flight management controller [13] as a low-level flight control unit including an inertial navigation estimator [14]. The ODROID is connected to the PX4 via a serial link and communicates with a ground-station via WiFi. It captures image data up to 24.6Hz with 1280x960 resolution from a forward-looking MatrixVision Bluefox2 camera connected via USB 2.0. The MAP uses the camera for sensing context-related information in the environment, enabling it to interact with the user (Fig. 1). All high-level tasks, including processing of image data, processing estimated poses from the motion tracker and sending feedforward correction data to the laser projection system, run onboard and are implemented in the ROS framework [15].

B. Laser Projection System

We custom-built a laser projector with a weight of approx. 100g from the following components: A 5mW laser module emits green light at a wavelength of 532nm. Due to the human eyes sensitivity to this wavelength, it improves contrast perception. The laser beam is deflected by a two-axis (tip/tilt) MEMS mirror [16]. Its reflective surface can be steered in a range of $\pm 5^\circ$ by applying a bias differential driving scheme. To generate the required DC bias voltage, we use an amplifier interfaced to a microcontroller (μC) via SPI, the SPI clock frequency is 1MHz. The μC is connected to the ODROID via a serial link running at 57.6Kb/s and relying on the ROS-serial package. To draw images with the laser, we use a vector graphics approach. For a demonstration

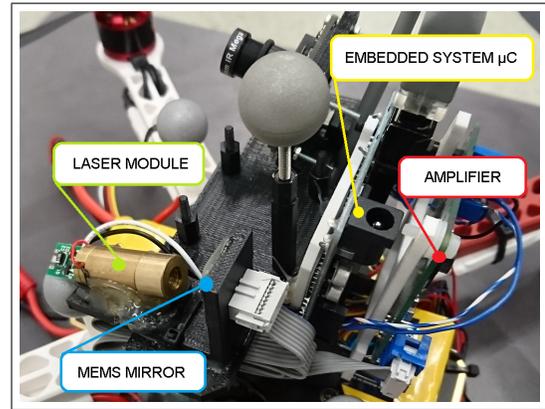


Fig. 4: Detailed overview of the laser projection system. The emitted light of the laser module is reflected by a forward facing MEMS mirror, interfaced to an amplification stage that converts commands from an embedded system (μC) into appropriate voltages and steers the mirror in two directions.

of projections, refer to Fig. 6. We send an array of ROS messages with x/y coordinates to the laser, whereby a rate of 100Hz is selected to improve connection reliability of the ROS serial link. The projector coordinates are defined by α_x and α_y and internally represented as angles at which the mirror is steered in x/y direction. Typical latencies of sending the messages from the SBC to the μC are 11ms. Sending position commands from the μC to the amplifier is currently done at a rate of 2kHz. As the complexity of measuring the time delay between sending commands from the μC and actually steering the MEMS mirror is significant, the delay is treated as unknown.

C. Flight Control Of The MAV

For flight control of the MAV, we use an Optitrack motion tracking system providing the PX4 flight management controller with low latency pose estimations. We use eight stationary cameras, covering an area of roughly $5 \times 4 \times 3m$ length, width and height. Poses from the Optitrack motion tracker are derived by the Optitrack server with a latency of 10ms. Transfer to the groundstation with Ethernet adds a latency below 1ms. To relate to measurement rates from a vision camera, poses are then delivered to the SBC via WiFi at 20Hz. Interfacing the PX4 flight management controller is also done via serial link running at 57.6Kb/s with below 10ms of delay. The system time between the Optitrack server, ground-station and the PX4 is synchronized via NTP/Chrony [17]. For a more detailed overview of the system, also including rates and latencies between the individual components, refer to Fig. 7.

D. Pose Estimation For Stabilization Of Projection

In addition to directly using the estimated poses from the motion tracking system, we further utilize an inertial position estimation approach on the low-level flight management controller to reduce noise and obtain interpolated position estimates at higher rates to improve stabilization

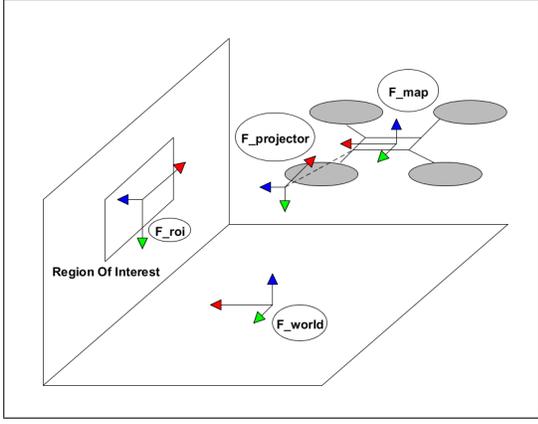


Fig. 5: Overview of coordinate frames in our experimental setup. Using right handed coordinate convention, the x-axis is colored red, y-axis is green and z-axis is blue.

accuracy of projections. We further exploit the noise filtering characteristics of the estimator to reduce flickering of the projected images and increase visual quality.

IV. PROJECTION STABILIZATION

In this section, we concentrate on the problem of stabilizing projected information suffering from movements of the MAV. We propose a method for stabilization and describe our implementation deployed on the onboard computer.

A. Coordinate Frames and Transformations

Fig. 5 shows the reference frames of our experimental setup: The world, the region of interest (ROI) on the wall, the MAP and the laser projector. Note that all coordinate frames are right handed, the coordinate frame of the projector and the ROI are in OpenCV conventions, letting the z-axis point towards the wall. We stabilize the projected image using the 4×4 homogeneous transformation matrix T_{PR} from the ROI on the wall with respect to the projector (Eq. 2), consisting of the rigid transformations T_{PM} from the MAP wrt. the projector, T_{MW} from the world wrt. the MAP and T_{WR} from ROI wrt. the world coordinate frame, which is tracked by the motion tracking system:

$$T_{PR} = T_{PM} \cdot T_{MW} \cdot T_{WR} \quad (1)$$

A 3D point P_R in the frame of the ROI can be transformed to the projector frame with

$$P_P = T_{PR}P_R \quad (2)$$

and further on into image space of the laser.

B. Laser Projector Model

As our projection system is based on a MEMS mirror, it shows nonlinear relations between applied DC voltage and mirror inclination of the individual axis [16]. To compensate for those nonlinearities, while still being able to relate to common computer vision algorithms, we suggest a pinhole camera model with nonlinear lens distortion to describe the characteristics of the projector. In our approach, we transform

a 3D point P_P given in the projector frame (Eq. 2) into the projectors coordinates.

By first normalizing the point P_P , we have

$$\begin{pmatrix} x' \\ y' \\ 1 \end{pmatrix} = \begin{pmatrix} P_{Px,n} \\ P_{Py,n} \\ 1 \end{pmatrix} = \begin{pmatrix} \frac{P_{Px}}{P_{Pz}} \\ \frac{P_{Py}}{P_{Pz}} \\ 1 \end{pmatrix} \quad (3)$$

Applying radial distortion with the Taylor approximation to an arbitrary function $L(r)$ [18] and neglecting all coefficients K_n except for the second-order term results in

$$\begin{pmatrix} x'' \\ y'' \end{pmatrix} = L(r) \begin{pmatrix} x' \\ y' \end{pmatrix} = \begin{pmatrix} (1 + K_{2x}r^2) \cdot x' \\ (1 + K_{2y}r^2) \cdot y' \end{pmatrix} \quad (4)$$

where x'' and y'' are the distorted coordinates, K_{2x} and K_{2y} are the distortion coefficients for the second-order terms and r is the radial distance $r^2 = x'^2 + y'^2$. We use two different distortion coefficients for the second-order terms of x and y axis, because we expect different nonlinear behaviors in terms of the individual axis of the MEMS mirror. This would be the equivalent to describing a distorted lens with an ellipsoid shape.

Finally, we apply calibrated intrinsics from the pinhole camera model, described in Section IV-B.1, to receive the actual coordinates in our laser image space.

$$\begin{aligned} \alpha_x &= f_x \cdot x'' + c_x \\ \alpha_y &= f_y \cdot y'' + c_y \end{aligned} \quad (5)$$

The projector coordinates α_x and α_y are directly used as the inputs for the laser projection interface to steer the mirror in x/y direction.

1) *Laser Model Calibration:* For the calibration of the model, we take a set of 6DOF poses of the MAP in the world coordinate frame. For each pose, we project a symmetric point grid pattern onto a planar wall and measure the 3D coordinates of the projected points in world coordinates. A desired grid of laser image points, gets distorted on a planar projection surface due to nonlinearities in the MEMS mirror. Compared to common camera calibration procedures, where a non-distorted checkerboard is used to calibrate for the intrinsics of the camera, our approach considers the opposite. The desired points in the laser space are the true grid points and get distorted in real world.

We use the Ceres Solver [19] and define the reprojection error in the laser image space as optimization residual. This has the advantage that the projection surface does not need to be known, as the calibration 3D projected points are directly measured. The natural choice would be to compare the 3D projected point positions from the calibration dataset against the projection of the laser commands into the real world. However, this would pose problems due to the non-existence of an analytic inverse for the camera distortion model, which needs to be coded into the residual.

Our solution is to utilize a reprojection error residual in the laser image space (Eq. 6). As a consequence, we need to propagate the measurement covariance into the laser image. This results in the forward covariance propagation

from the measured 3D point to the laser image through the backprojection operation. Thus, our calibration respects the maximum likelihood principle and can be used to estimate a proper calibration parameter covariance through the backward transport of covariance theorem [18].

Our calibrated parameters are defined by

$$\mathbf{r}(\mathbf{T}_{PR}, \mathbf{K}, \mathbf{D}) = \frac{\alpha_{img,cmd} - \hat{\alpha}_{img,backpr.}}{\sigma} \quad (6)$$

where T_{PR} transforms a 3D point measured on the wall into the laser projector frame, K and D include the calibrated intrinsic parameters and distortion coefficients of the laser model, $\alpha_{img,cmd}$ are the commanded points in the projector image plane, $\hat{\alpha}_{img,backpr.}$ are the points in the projector image plane derived from the backprojection and σ is the previously discussed propagated covariance of our 3D points, measured with the motion tracker, through the backprojection operation.

C. Compensation via Feedforward Correction

For stabilization of the projected information, we want to use a simple but robust feedforward correction algorithm. We define our desired projection in 3D, which is represented by a projected point on a wall. The position and shape of the target projection surface are known and used to calculate P_R . The desired 3D point P_R is defined in the ROI frame F_{roi} (please also refer to Fig. 5 and Fig. 10a). During our experiments, based on this desired point, we derive our projector coordinates and steer the mirror towards it. We do this in the following way: We transform P_R into coordinate frame of the MAP F_{map} using $T_{MW} \cdot T_{WR}$ based on poses from the motion tracker. Next, we use the transformation T_{PM} , calibrated from MAP to projector, to derive coordinates of the defined point in the laser projector frame $F_{proj.}$. Including the intrinsics and radial distortion from our projector model, we finally steer the mirror to the position of the 3D point P_R . According to Fig. 9, while we calculate the coordinates for correction on the SBC and forward commands to the laser interface (time delays t_1 to t_3), the MAP is moving. Due to the delays of the laser interface, the actual projection happens when the MAP has moved already ($t_3 + \Delta t$). Thus, we are not able to compensate for the exact position of the desired 3D point P_R , which results in an offset of the projection. This problem affects all compensation methods used in our experiments.

For an improved stabilization approach, including noise filtering and interpolation between rare pose readings, we utilize the inertial state estimator of the onboard flight management controller (Fig. 7). The poses from the motion tracking system are sent at a rate of 20Hz to the SBC and arrive with approximately 21ms of delay. The inertial state estimator fuses measurements of the IMU, running at high rates, with poses from the motion tracker. The state estimator predicts and corrects position in x, y and z for the current time step, whereby predictions are based on accelerometer measurements. It filters out noise and interpolates the motion tracker poses, which are derived from the flight management

controller at a rate of 100Hz. This estimated pose is again used to transform the coordinates of the 3D point P_R into the frame of the laser projector. For correction, the updated coordinates in laser space are finally forwarded to the laser projection interface via the serial link. The latency is thereby approximately 11ms.

Fig. 8 shows a timing diagram between the individual system components with intermediate rates and latencies.

V. EXPERIMENTS

In this section, we analyze the performance of the stabilized projection during flight. We report on four experiments: We elaborate on the accuracy of feedforward stabilization based on poses directly derived from Optitrack and compare the results with compensation by using the fused poses from the inertial estimator of the low level flight management controller. Both methods are applied during hover flight and also during dynamic flight. An example for the MAP's actual position trajectories is shown in Fig. 10b.

A. Hover Flight

For experiments during hover flight, we want to position the MAP in air so that the 3D point P_R can be "seen" by the FOV of the laser projector. At the beginning of the experiment, we place the MAP at the origin of the world frame facing the wall. When the MAP takes off, it is commanded to a hovering height of 0.65m, again to be able to project close to the 3D point. The MAP's attitude setpoint is oriented towards the ROI. As we can localize the pose of the ROI on the wall with the motion tracking system, we use our onboard feedforward correction algorithm to project the desired point into the origin of the ROI. At the same time, we project the uncompensated point $\alpha_x/\alpha_y = (0, 0)$ in projector coordinates which represents the disturbance due to the movements of the MAV. To quantify the effects of the compensation, we capture image data of the ROI with an external camera over 60s with a resolution of 1920x1080 at 25fps and detect the uncompensated and compensated points in the image. The ROI is thereby of size A3 with WxH of 420x297mm. The 1500 frames are downsampled by a rate of 10 which results in 150 point pairs (compensated and uncompensated points) in every experiment. In every image, we use the Euclidean norm of x/y coordinates in F_{roi} and measure the distance of the projected uncompensated points and compensated points to the desired 3D point P_R . We derive mean error and standard deviation (\bar{X} and σ).

B. Dynamic Flight (Circle Flight)

For dynamic flight we calculate a full 3D elliptical position trajectory with the center located at $x/y = (0, 0)m$ in the world frame and 0.65m above ground. In the x/y plane, we command a circular trajectory with a diameter of 1.4m and with an angular speed of $36^\circ/s$. During flight, we set the attitude of the MAP towards the walls ROI and adjust the height to keep the 3D point P_R approximately in the laser projectors FOV. The wall is located at a distance of 2.18m towards positive x-direction in the world frame.

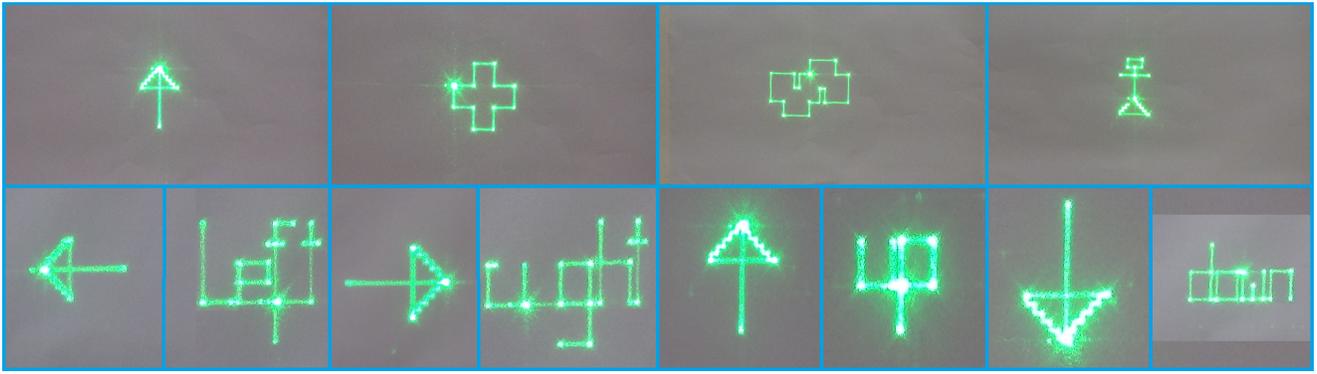


Fig. 6: Shown is a set of symbols which can be projected by the projection system. Included are an arrow for instructing users, a red cross representative for a point of interest and the logo of the Graz University of Technology. Additionally we can project directions and letters, for example, to guide a user inside of a building.

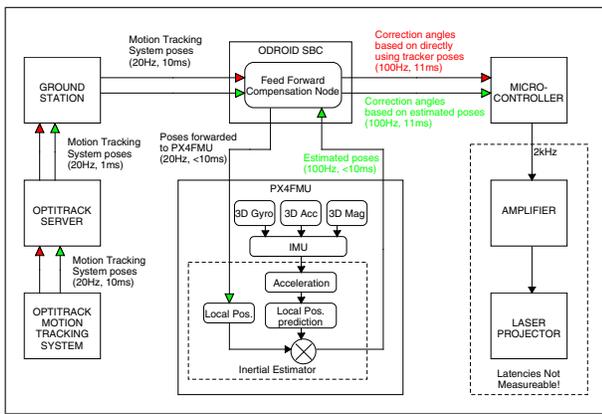


Fig. 7: Overview of rates and latencies of the experimental setup. The red path shows the poses which are transferred in terms of the direct compensation approach. The green path indicates transfer of poses for compensation via the inertial estimator.

TABLE I: Error characteristics of feedforward compensation.

Mean(Std.Dev.) [mm]	uncomp.	comp.
$\bar{X}(\sigma)_{H,noinav}$	23.5 (38.7)	5.2 (6.6)
$\bar{X}(\sigma)_{H,inav}$	27.7 (40.3)	8.6 (5.0)
$\bar{X}(\sigma)_{C,noinav}$	47.5 (52.8)	7.4 (21.0)
$\bar{X}(\sigma)_{C,inav}$	66.6 (56.0)	7.3 (19.3)
$\bar{X}(\sigma)_{S,noinav}$	221.9 (54.9)	7.5 (8.4)
$\bar{X}(\sigma)_{S,inav}$	203.9 (47.4)	9.8 (5.8)

C. Experimental Results

Table I shows a summary of the four experiments. We provide mean and standard deviation of the uncompensated projected points compared to the compensated ones during hover flight $\bar{X}(\sigma)_H$ and circle flight $\bar{X}(\sigma)_C$. Both experiments are evaluated either with or without using inertial estimates.

In Fig. 11a, uncompensated and compensated projected points during the circle flight are shown. The results are

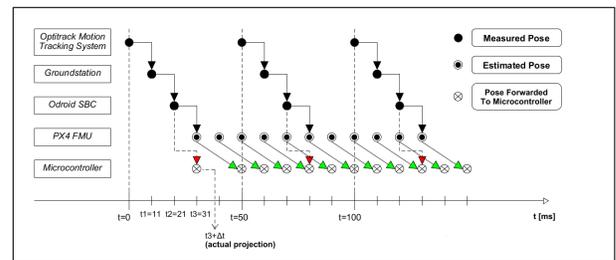


Fig. 8: Timing of the poses used for feedforward compensation. The red arrows indicate poses of the motion tracking system directly used for compensation, whereas green arrows are poses derived from the inertial estimator.

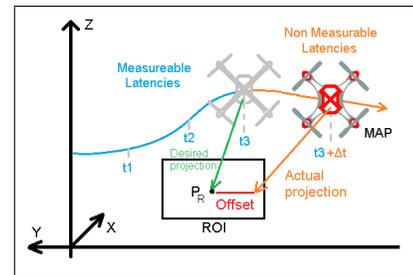
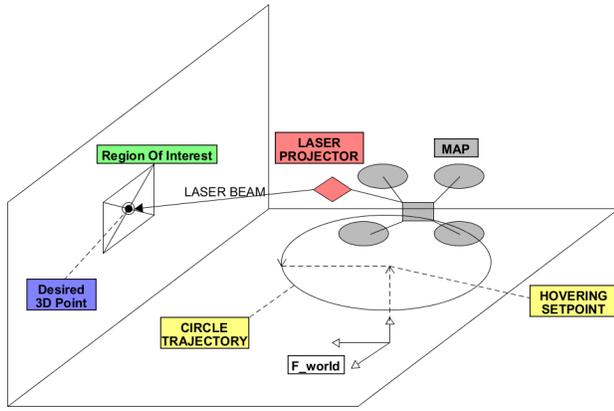


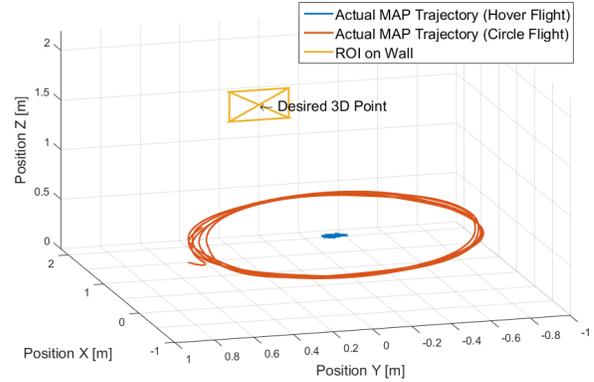
Fig. 9: Effect of the laser interfaces unknown latency. After the coordinates for correction are derived and forwarded to the laser projector, any movement of the drone results in an offset of the projection.

based on directly using the pose estimates from the motion tracking system. It is clearly visible that the dynamic error, resulting from the movements of the drone, is significantly reduced by the feedforward correction. Offsets to the desired point in the origin of the axis are a result of movements of the MAP, errors in the calibration of the transformation T_{PM} and inaccuracies of the estimated model of the laser projector.

Fig. 11b shows results of the same experiment, but the feedforward compensation is based on poses from the inertial estimator. It is obvious that the filtering characteristics of the inertial estimator reduce noise and also help to smooth out



(a) Sketch of the experimental setup.



(b) Actual position trajectories.

Fig. 10: The experiments of the MAP in detail. (a) A sketch of the experimental setup. Fig. (b) The actual position trajectories of the MAP during circle and hover flight based on a view from behind.

the movements of the projected image on the plane. This improves accuracy, which is also reflected in the standard deviations $\sigma_{C,noinav}$ and $\sigma_{C,inav}$. Moreover, also visual quality is significantly increased. Therefore, please refer to the supplementary video.

We took a set of measurements positioning the MAP at random static poses in front of the ROI, to inspect on the static error of the feedforward compensation. The uncompensated point is steered randomly to outer regions of the FOV, which also means that the MEMS mirror almost reaches its maximum inclination angles. It reflects accuracy of the calibration and influences of nonlinearities of the laser model.

It is noticeable that the error characteristics of compensated points $\sigma_{H,noinav}$ and $\sigma_{H,inav}$ during hover flight are slightly below the static error characteristics $\sigma_{S,noinav}$ and $\sigma_{S,inav}$. This is due to the fact that, in hovering position, the dynamics of the MAP are low and the origin of the laser projector can be kept quite close to the desired 3D point. Therefore, inaccuracies in our estimated nonlinear laser model have less influence on the compensated points. Furthermore, the mean error during hover flight without using estimates ($\bar{X}_{H,noinav}$) is significantly lower. This is considered to be a result of the low disturbance characteristics during this specific flight experiment.

D. Use Case Scenario

As we think of the MAP as a robotic companion, it should be able to support the inexperienced user in certain HMI scenarios. Therefore we implemented a "teaching assistant" scenario as a showcase. The MAP is able to support a student with solving basic equations (Fig. 1) by autonomously detecting the content of the equation, solving it and projecting the result next to it. Please also refer to the supplementary video.¹

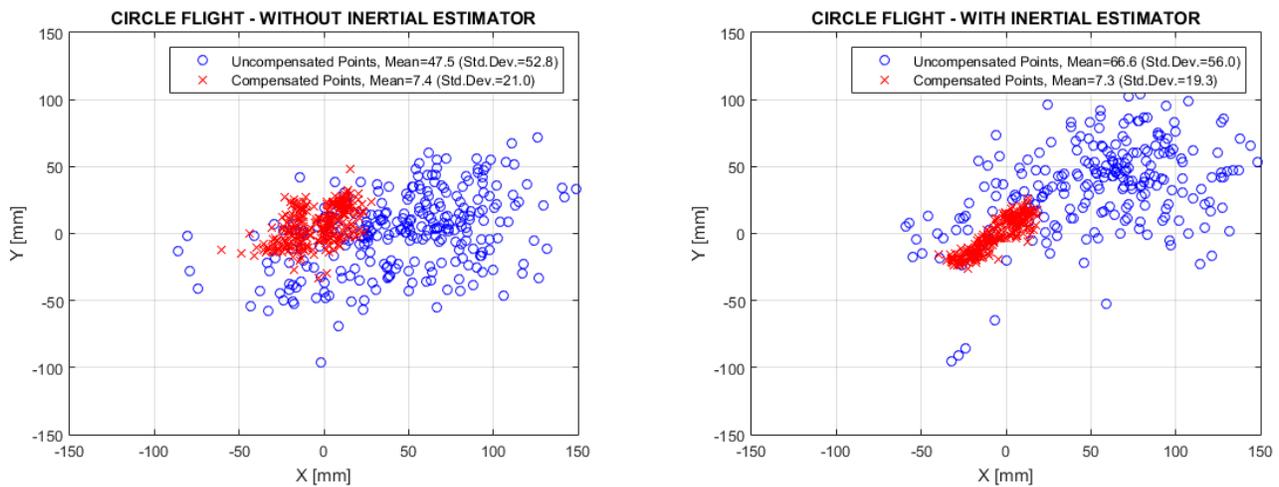
¹Supplementary video: <https://youtu.be/0A2EtqAMNE>

In detail, the following scenario is shown: The MAP first approaches a given equation and reads it during flight by utilizing a neural network based approach for text recognition. Then, the user approaches the equation and tries to solve it, while the MAP moves away to keep distance. In our example, the user fails on the first attempt. When the user writes down a result and moves away from the equation, the MAP comes in to check the result. Depending on validity of the result it projects "✓" for correct or "X" for wrong on the wall to signal the user validity of the result. Then it moves away again. The user should write down the correct result and move away from the equation. The MAP moves in again to check and signals if correct result was detected by projecting on the wall. Finally, the drone approaches a default waypoint and projects the correct result next to the equation. During the waiting phases, the MAP signals that the user should interact with the scene by projecting "hold" on the wall.

VI. DISCUSSION AND CONCLUSIONS

Evaluation of the MAP as our first step towards mixing airborne robotics and spatial augmented reality unveils a novel platform capable of projecting stabilized images in the 3D environment. It is able to stabilize the projection during hover and even dynamic flight and increases visual quality by utilizing the onboard hardware. Combined with compact size and acceptable flight times, the MAP is able to guide and support the user in various scenarios.

Limiting factors are the narrow field of view of 10° of the laser projector and current dependency on the motion tracking system. We are able to overcome limitation, related to the FOV, by utilizing the dynamics of the MAP to steer the laser approximately towards the desired projection region. In addition, we introduce a laser projection model, which is able to compensate for nonlinearities. With mean projection error characteristics of below 1cm at a maximum distance of approximately 2.9m, we are able to compensate



(a) Results of compensation based on directly using Motion Tracker.

(b) Results of compensation using estimated poses.

Fig. 11: Measured projections of uncompensated and compensated points during circle flight. (a) Results of compensation based on directly using poses from the Optitrack motion tracking system. (b) Results of compensation using estimated poses from the flight management controller.

for higher dynamics of the system and improve visual quality of projected symbols.

As a future work we want to implement state prediction considering the time delays of the system to predict the movement of the MAP to better compensate the latency of the laser projector interface. As a next step, the position and shape of the target projection surface could be estimated online in combination with the onboard vision camera, as these are currently known variables used to calculate P_R . With such an extension, the MAP could dynamically adapt to non-planar, skewed or tilted projection surfaces, and even adapt the projected symbol to the current user viewpoint. We also foresee using SLAM or optical flow for navigation to improve mobile capabilities and make the platform more flexible for future HMI scenarios.

REFERENCES

- [1] Oliver Bimber and Ramesh Raskar. *Spatial Augmented Reality: Merging Real and Virtual Worlds*. A. K. Peters, Ltd., Natick, MA, USA, 2005.
- [2] Claudio S. Pinhanez. The everywhere displays projector: A device to create ubiquitous graphical interfaces. In *Proceedings of the 3rd International Conference on Ubiquitous Computing, UbiComp '01*, pages 315–331, London, UK, UK, 2001. Springer-Verlag.
- [3] Andrew Wilson, Hrvoje Benko, Shahram Izadi, and Otmar Hilliges. Steerable augmented reality with the beamatron. In *Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology, UIST '12*, pages 413–422, New York, NY, USA, 2012. ACM.
- [4] Horst Hörtnner, Matthew Gardiner, Roland Haring, Christopher Lindinger, and Florian Berger. Spaxels, pixels in space—a novel mode of spatial display. In *SIGMAP*, pages 19–24, 2012.
- [5] Chris Harrison, Hrvoje Benko, and Andrew D. Wilson. Omnitouch: Wearable multitouch interaction everywhere. In *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology, UIST '11*, pages 441–450, New York, NY, USA, 2011. ACM.
- [6] Keisuke Tajimi, Keiji Uemura, Yasuhiro Kajiwara, Nobuchika Sakata, and Shogo Nishida. Stabilization method for a hip-mounted projector using an inertial sensor. In *20th International Conference on Artificial Reality and Telexistence (ICAT)*, Adelaide, Australia, 2010. IEEE Computer Society.
- [7] N. Sakata, T. Kurata, T. Kato, M. Kourogi, and H. Kuzuoka. WacI: supporting telecommunications using - wearable active camera with laser pointer. In *Wearable Computers, 2003. Proceedings. Seventh IEEE International Symposium on*, pages 53–56, Oct 2003.
- [8] Visual assist with a laser pointer and wearable display for remote collaboration. In *Proc. 2nd International Conference on Collaboration Technologies*, pages 66–71, 2006.
- [9] T. Maeda and H. Ando. Wearable scanning laser projector (wslp) for augmenting shared space. In *ACM SIGGRAPH 2004 Sketches, SIGGRAPH '04*, pages 109–, New York, NY, USA, 2004. ACM.
- [10] Björn Schwerdtfeger, Daniel Pustka, Andreas Hofhauser, and Gudrun Klinker. Using laser projectors for augmented reality. In *Proceedings of the 2008 ACM Symposium on Virtual Reality Software and Technology, VRST '08*, pages 134–137, New York, NY, USA, 2008. ACM.
- [11] Jürgen Scheible, Achim Hoth, Julian Saal, and Haifeng Su. Display-drone: A flying robot based interactive display. In *Proceedings of the 2nd ACM International Symposium on Pervasive Displays, PerDis '13*, pages 49–54, New York, NY, USA, 2013. ACM.
- [12] Y. Hosomizo, D. Iwai, and K. Sato. A flying projector stabilizing image fluctuation. In *Consumer Electronics (GCCE), 2014 IEEE 3rd Global Conference on*, pages 31–32, Oct 2014.
- [13] Lorenz Meier, Petri Tanskanen, Friedrich Fraundorfer, and Marc Pollefeys. Pixhawk: A system for autonomous flight using onboard computer vision. In *Robotics and Automation (ICRA), 2011 IEEE International Conference on*, pages 2992–2997, 2011.
- [14] Paul Riseborough. Inertialnav. <https://github.com/priseborough/InertialNav>, 2015.
- [15] Morgan Quigley, Ken Conley, Brian P. Gerkey, Josh Faust, Tully Foote, Jeremy Leibs, Rob Wheeler, and Andrew Y. Ng. Ros: an open-source robot operating system. In *ICRA Workshop on Open Source Software*, 2009.
- [16] V. Milanovic, G.A. Matus, and D.T. McCormick. Gimbal-less monolithic silicon actuators for tip-tilt-piston micromirror applications. *Selected Topics in Quantum Electronics, IEEE Journal of*, 10(3):462–471, May 2004.
- [17] User guide for the chrony suite. <https://web.archive.org/web/20150907232505/http://chrony.tuxfamily.org/manual.html>. Accessed: 2015-09-14.
- [18] R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, ISBN: 0521540518, second edition, 2004.
- [19] Sameer Agarwal, Keir Mierle, and Others. Ceres solver. <http://ceres-solver.org>.