

©2019 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Please cite this paper as:

```
@InProceedings{Hinzmann2019_trinocular,
  Title          = {Flexible Trinocular: Non-rigid Multi-Camera-IMU Dense
                    Reconstruction for UAV Navigation and Mapping},
  Author         = {Hinzmann, Timo and Cadena, Cesar and Nieto, Juan and
                    Siegwart, Roland},
  Booktitle      = {2019 {IEEE/RSJ} International Conference on
                    Intelligent Robots and Systems ({IROS})},
  Year           = {2019},
  Organization   = {IEEE}
}
```

Flexible Trinocular: Non-rigid Multi-Camera-IMU Dense Reconstruction for UAV Navigation and Mapping

Timo Hinzmann, Cesar Cadena, Juan Nieto, and Roland Siegwart

Abstract—In this paper, we propose a visual-inertial framework able to efficiently estimate the camera poses of a non-rigid trinocular baseline for long-range depth estimation on-board a fast moving aerial platform. The estimation of the time-varying baseline is based on relative inertial measurements, a photometric relative pose optimizer, and a probabilistic wing model fused in an efficient Extended Kalman Filter (EKF) formulation. The estimated depth measurements can be integrated into a geo-referenced global map to render a reconstruction of the environment useful for local replanning algorithms. Based on extensive real-world experiments we describe the challenges and solutions for obtaining the probabilistic wing model, reliable relative inertial measurements, and vision-based relative pose updates and demonstrate the computational efficiency and robustness of the overall system under challenging conditions.

I. INTRODUCTION

For environmental awareness and autonomous mission execution in previously unknown terrain unmanned aerial platforms depend on high-quality depth measurements and ideally have access to a globally consistent map. For ground robots and rotary-wing Unmanned Aerial Vehicles (UAVs), small baseline rigid stereo rigs or depth cameras have shown to be sufficient to enable autonomous operation. However, for small fixed-wing UAVs flying at 15 m/s, like the platform shown in Fig. 1, off-the-shelf sensor solutions fail for various reasons: Monocular depth from motion (high depth uncertainty due to epipole), small baseline rigid stereo (limited range and depth uncertainty), laser (limited range, heavy, and expensive), depth cameras (limited range), radar (not yet miniaturized).

In this paper, we present our approach to increase the environmental awareness of fixed-wing UAVs to efficiently estimate the time-varying relative pose between multiple cameras while only relying on inexpensive visual-inertial sensors. In contrast to [1], in this paper we propose a trinocular setup where the center camera is rigidly mounted with respect to the autopilot located inside the fuselage to simplify the generation of a geo-referenced and globally consistent map. Additionally, the trinocular setup allows to take advantage of both the full baseline for long-range depth estimates (left to right wing tip camera; approx. 2.4 m for our platform shown in Fig. 1) and half the baseline (e.g. left to center camera) for obstacles nearby or for landing procedures. We estimate the non-rigid time-varying baseline between the

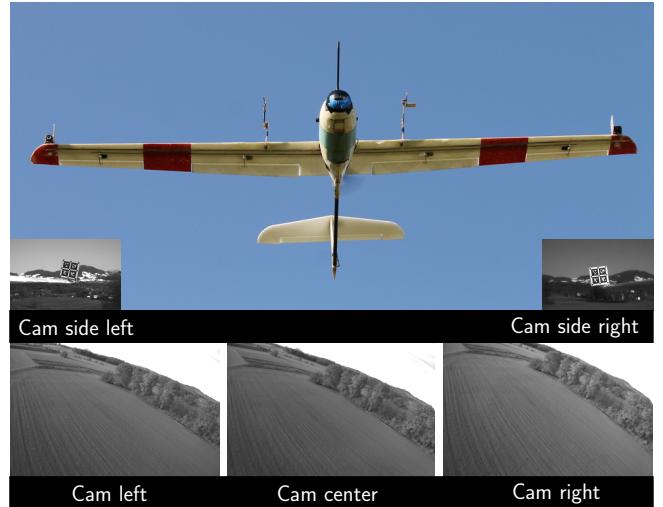


Fig. 1: Fixed-wing UAV platform *Techpod* equipped with the proposed trinocular visual-inertial sensor setup. The side cameras and April tags are used for the identification of the wing model.

center and left/right camera with Extended Kalman Filters (EKFs) that fuse relative visual-inertial measurements with a calibrated wing model that further constrains the relative baseline transformation. The depth maps resulting from the multiple stereo views can then be integrated into a geo-referenced map. While the approach presented in [1] is interesting for reactive navigation and control of a UAV, the framework proposed in this paper is designed for the use of model predictive controllers and local replanning algorithms able to operate on a 3D or 2.5D map.

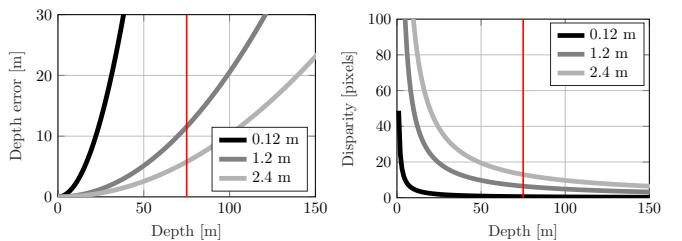


Fig. 2: Theoretical depth error and disparity for different stereo baselines in meters. The red vertical line marks the distance that our UAV platform can cover within 5 s. Further parameters: focal length of 2.8 mm, pixel size of 6 μ m, assumed resolvable disparity of one pixel.

*This work was supported by the Federal office armasuisse Science and Technology under project n°050-45.

All authors are with the ETH, the Swiss Federal Institute of Technology Zurich, Autonomous Systems Lab (www.as.ethz.ch), Leonhardstrasse 21, LEE, CH-8092 Zurich, Switzerland. {firstname.lastname}@mavt.ethz.ch

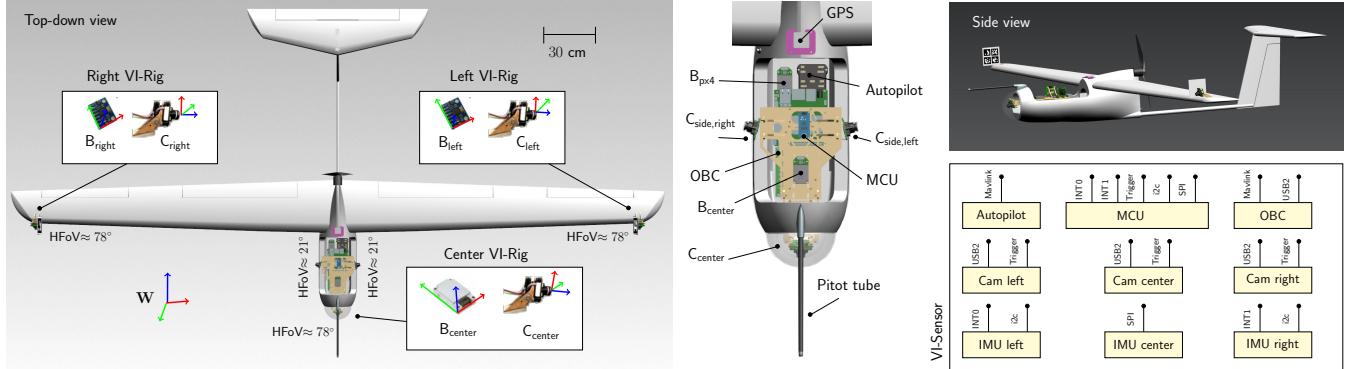


Fig. 3: Overview showing the UAV platform, sensors, and coordinate frames. An EKF estimates the relative pose between left, right, and center camera. The side cameras are used to identify a wing model in form of a Gaussian prior. The horizontal field of view (HFoV) is denoted for each camera. The schematic on the right bottom depicts the developed VI-sensor that time-synchronizes the sensors on a hardware level.

II. RELATED WORK

Several state-of-the-art monocular visual-inertial odometry (VIO) and simultaneous localization and mapping (SLAM) frameworks have been adopted to stereo or multi-camera operation [2]–[6]. Other multi-camera systems are specifically designed for their individual application such as for increased environmental awareness to navigate in cluttered environments using a multi-copter [7], [8] or for obstacle detection and avoidance on-board an agile delta wing [9]. All of these approaches have in common that they assume a fixed and precisely calibrated transformation between the individual cameras. However, for a fixed-wing UAV as shown in Fig. 1, a *rigid* stereo system can, due to aerodynamic effects and payload constraints, only be mounted inside the fuselage, limiting the effective baseline in this case to approximately 12 cm. Fig. 2 presents the theoretical depth uncertainty and expected disparity for such a rigid small-baseline stereo setup and motivates the necessity for a setup with a wider baseline.

On the other hand, approaches have been developed that are able to cope with different degrees of deviations from the calibrated baseline transformation: Warren et al. [10] continuously estimate thermally induced slow changes in the baseline. In [11], a down-looking stereo pair with a wide-baseline (0.7 m) is employed and the relative transform between the cameras, together with the poses of the stereo rig itself is estimated in an offline bundle adjustment problem. Lanier et al. [12]–[14] use a set of accelerometers distributed along the wing to estimate vibrational disturbances without any vision update. For airborne applications, Yang et al. fuse the measurements from a master and slave IMU in an EKF to estimate the relative states, a process referred to as transfer alignment [15]. Achtelik et al. [16] proposed an efficient EKF formulation that fuses relative IMU and vision measurements for stereo vision with two quadrocopters that have an overlapping field of view. Recently, this EKF formulation has been used in [1], [17], [18] for various applications. In [1], the initial concept of *Flexible Stereo* was introduced,

further constraining the EKF formulation with a probabilistic wing model. The approach allows obstacle avoidance with a reactive controller as the depth map is computed with respect to the left (or right) wing tip camera.

In contrast to [1], this paper suggests to use a trinocular setup to enable the generation of a geo-referenced and globally consistent map. Furthermore, we suggest a vision update that utilizes photometric feature tracking in combination with relative inertial data to obtain an improved initial position. Whilst [1] is only tested in simulation, this publication demonstrates the real-world application.

III. CONTRIBUTIONS AND ASSUMPTIONS

This paper incorporates recent advances in relative visual-inertial state estimation [1], [19], [20] and contains the following contributions:

- A procedure to obtain a probabilistic wing model that considers the aerodynamic forces acting on the wing during the flight.
- Implementation and validation of an efficient framework to estimate the time-varying relative transformation between multiple stereo setups.
- Real world experiments, including platform design and hardware considerations.

To focus on the description of the above-mentioned contributions we assume that the center camera is aligned with the map throughout the paper. For flying in GPS-denied environments this can be realized for instance with [21] or [20]. As the state estimator implemented on our autopilot already provides pose priors in real-time we instead employ the approach suggested in [22] but augment it with IMU edges. In summary, the proposed framework assumes the following input: 1) features with corresponding depth estimates tracked in the center camera for relative camera-camera alignment, 2) pose estimates of the center camera in the map frame for depth map registration.

IV. THE APPROACH

In Sec. IV-A we present our approach to align the wing camera's pose to the center camera. All steps in this section follow the paradigm of staying close to the optimal relative pose to increase the efficiency of the photometric image alignment and relative pose estimation. Sec. IV-B shortly summarizes the generation of the depth maps while Sec. IV-C describes our procedure to obtain a probabilistic wing model.

A. Inertial-photometric alignment of wing to center camera

The pose of a wing camera is aligned with the center camera using an EKF in relative formulation [19], efficiently fusing relative IMU measurements, photometric visual cues, and a probabilistic wing model [1]. The state vector is defined as follows

$$\mathbf{x} = [\bar{q}_c^j, \omega_c^c, \omega_j^j, \mathbf{p}_c^j, {}_w\mathbf{v}_c^j, {}_w\mathbf{a}_c^c, {}_w\mathbf{a}_j^j, \mathbf{b}_{\omega_c}^j, \mathbf{b}_{\omega_c}^j]^T \quad (1)$$

with $j \in \{\text{left}, \text{right}\}$. In comparison to [1], [19], the state is augmented to compensate for the IMU bias. Formulating the relative pose alignment in a relative EKF is efficient and modular, allowing to easily add additional visual-inertial sensor rigs to, e.g., further increase the overall field of view.

Based on the last iteration's estimate of $\mathbf{T}_{C_{c,k-1}}^{C_{j,k-1}}$, the relative transformation is propagated to the time stamp of the current stereo frame using the IMU measurements of the center and wing camera (as outlined in [1], [19]), resulting in an initial prior for $\mathbf{T}_{C_{c,k}}^{C_{j,k}}$ at time step k . This relative pose prior is then used to project features observed in the center image into the wing camera via

$$\begin{aligned} d &= \|\mathbf{p}_{lm}^W - \mathbf{p}_{C_{c,k}}^W\|_2 \\ \mathbf{u}_j &= \Pi(\mathbf{T}_{C_{c,k}}^{C_{j,k}} \cdot \mathbf{C}_c \cdot d) \end{aligned} \quad (2)$$

where \mathbf{p}_{lm}^W is the estimated landmark position, d is the observed sparse depth estimate corresponding to this landmark, $\Pi(\cdot)$ is the projection operator, and \mathbf{C}_c is the bearing vector from the center camera to the landmark.

Patches around the expected set of feature positions \mathbf{u}_j are extracted from the wing camera image and the relative transformation aligning the wing to the center camera is further refined using sparse image alignment. The photometric error is minimized as proposed in [20], [23], using a constrained Gauss-Newton solver with pyramidal layers:

$$\begin{aligned} \mathbf{T}_{C_{c,k}}^{C_{j,k}} &= \arg \min_{\mathbf{T}_{C_{c,k}}^{C_{j,k}}} \frac{1}{2} \sum_i \|\delta \mathbf{I}(\mathbf{T}_{C_{c,k}}^{C_{j,k}}, \mathbf{u}_{j,i})\|_{\Sigma_I}^2 \\ &+ \frac{1}{2} \|\mathbf{p}_{C_{c,k}}^{C_{j,k}} - \bar{\mathbf{p}}_{C_c}^{C_j}\|_{\Sigma_p} + \frac{1}{2} \|\log(\bar{\mathbf{R}}_{C_c}^{C_j \top} \mathbf{R}_{C_{c,k}}^{C_{j,k}})^{\vee}\|_{\Sigma_R}. \end{aligned} \quad (3)$$

For the derivation and notation details we refer to [20], [23]. As motion constraint we choose the transformation prior $\bar{\mathbf{T}}_{C_c}^{C_j} = (\bar{\mathbf{R}}_{C_c}^{C_j}, \bar{\mathbf{p}}_{C_c}^{C_j})$ obtained from the wing model. Alternatively, the current mean and covariance from the EKF could be used. While the translation computed in the vision update presented in [1], [19] is only up to scale, the approach described in this publication computes a full-scale translation. The relative pose from the vision step is then

fused in the EKF with the Gaussian wing model as in [1], for the case that the vision step fails, and the framework continues with the next iteration.

a) Discussion: In our experiments, above approach has shown to be a robust way to transfer and track features from one to another camera under challenging conditions (e.g. different illumination or exposure time) as the feature locations are constrained by the relative transformation and are not tracked individually. On the other hand, tracking each feature patch individually, for instance with a Lucas-Kanade (LK) tracker, was not robust and resulted in many outlier matches. Note that this LK tracker implementation showed a good performance for establishing time-sequential frame-to-frame correspondences for the *same* camera.

B. Dense reconstruction

The optimized camera poses are used to rectify the images using the algorithm proposed in [24] with subsequent stereo block-matching (BM). Based on the three cameras, the following combinations of depth maps can be computed: $\mathbf{d}_{c,l}^c$, $\mathbf{d}_{c,r}^c$, and $\mathbf{d}_{l,r}^c$ with their different characteristics visualized in Fig. 2. For instance $\mathbf{d}_{c,l}^c$ denotes the depth map computed from the left and center image, expressed in the center camera.

C. Obtaining a wing model

The probabilistic wing model used in this paper is represented by a Gaussian mean and covariance for the relative transformation between wing tip (left, right) and center, fuselage IMU. In order to capture the aerodynamic forces acting on the wing the relative transformation needs to be measured *in-air* as shown by Fig. 4: The pictures are recorded by the right side camera mounted rigidly inside the fuselage looking onto the wing towards the right wing camera (with attached April tag) and illustrate the aerodynamic lift force acting on the wing.

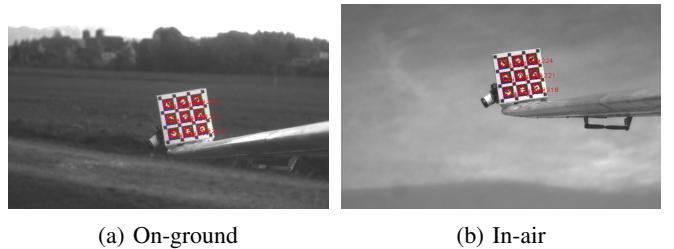


Fig. 4: Side view from camera rigidly mounted inside the fuselage, looking onto the right wing tip camera.

Our proposed calibration process consists of estimating

- 1) the rigid transformation $\mathbf{T}_{tag_j}^{C_j}$ (in the lab under static conditions) and
- 2) the time-varying transformation $\mathbf{T}_{tag_j}^{C_{side,j}}$ (using data recorded during a flight).

Firstly, we describe how to obtain the rigid transformation $\mathbf{T}_{tag_j}^{C_j}$ between the wing camera and the wing tag: The experiment setup and involved coordinate frames are shown

in Fig. 5. The rigid relative transformation between the wing camera and wing tag is given by

$$\mathbf{T}_{\text{tag}_j}^{C_j} = \mathbf{T}_{C_c}^{C_j} \mathbf{T}_{C_{\text{side},j}}^{C_c} \mathbf{T}_{\text{tag}_j}^{C_{\text{side},j}} \quad (4)$$

where the rigid transformations $\mathbf{T}_{C_c}^{C_j}$ and $\mathbf{T}_{C_{\text{side},j}}^{C_c} = \mathbf{T}_{B_c}^{C_c} \mathbf{T}_{C_{\text{side},j}}^{B_c}$ are calibrated with *Kalibr* [25].

Note that $\mathbf{T}_{\text{tag}_j}^{C_j}$ could also be obtained by solving the hand-eye problem (e.g. with [26]) but requires access to an accurate motion capture system (e.g. Vicon). Instead, our procedure relies only on the synchronized images and inertial measurements from the VI-sensor.

Secondly, during the flight, the April tags on the wing cameras are detected by the left respectively right side camera. Based on the April tag detection, the relative transformation $\mathbf{T}_{\text{tag}_j}^{C_{\text{side},j}}$ is obtained from an absolute pose estimator (PNP). That is, for every frame, the camera pose of the wing camera with respect to the center camera can be computed as

$$\mathbf{T}_{C_c}^{C_j} = \mathbf{T}_{\text{tag}_j}^{C_j} \mathbf{T}_{C_{\text{side},j}}^{\text{tag}_j} \mathbf{T}_{C_c}^{C_{\text{side},j}}. \quad (5)$$

Note that the side cameras and April tags need to be mounted only once for every UAV type in order to establish the wing model.

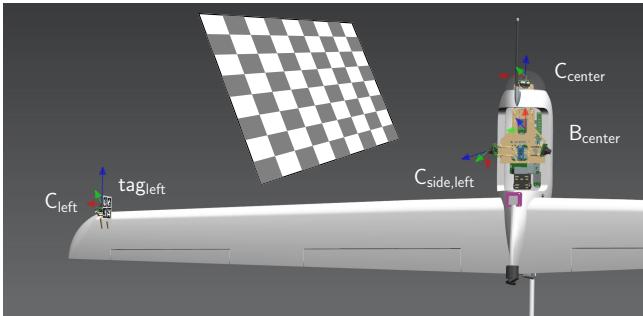


Fig. 5: Coordinate frames involved in the calibration procedure to obtain a probabilistic wing model (left side).

V. PLATFORM AND HARDWARE

Fig. 3 gives an overview of the employed UAV platform *Techpod* and of the VI-sensor system integration. Due to specific hardware requirements, such as the long baseline between the wing and center camera, commercially available VI-sensors [27] could not be used and a customized VI-sensor based on the *Atmega328P* microcontroller unit (MCU) was developed: The Camera-IMU rig for the left and right wing each consist of a 0.3 MP *Aptina MT9V034* camera and an IMU of type *MPU6050*, rigidly mounted behind the camera. The IMU measurements are triggered via interrupts and transferred to the MCU. Due to the required length of the data lines, three i^2C repeaters are required: one next to the MCU and two next to the Camera-IMU rigs on the wing just before the signal enters the *MPU6050*. The spatially closer and more accurate center IMU (*ADIS16448*) is connected to the MCU via SPI. The three IMUs and the five cameras of same type (three in normal mode, two additional for

the wing calibration) are time-synchronized via hardware lines and run at 100 Hz and up to 20 Hz respectively. The rigid, non-time-varying transformations $\mathbf{T}_{B_{\text{center}}}^{B_{\text{px4}}}$ and $\mathbf{T}_{C_j}^{B_j}$ with $j \in \{\text{left, right, center}\}$ are calibrated offline with *Kalibr* [25], where B_{px4} refers to the IMU (*ADIS16448*) used by the *pixhawk* Autopilot. An *UP Squared* with *Intel Atom* CPU at 1.6 GHz was used as on-board computer (OBC).

VI. EXPERIMENTS

A. Inertial measurements

In particular for the inertial measurement units mounted on the wing tips strong vibrations are to be expected. To quantify this effect, we compare *measured* to *expected* IMU readings. The camera poses $\mathbf{T}_{C_j}^W$ are optimized offline (using GPS and vision measurements only) and transformed to the IMU poses $\mathbf{T}_{B_j}^W$ via the rigid transformation $\mathbf{T}_{B_j}^{C_j}$. Using our implementation of [28] the expected linear accelerations and angular velocities are computed. The expected and measured (raw,

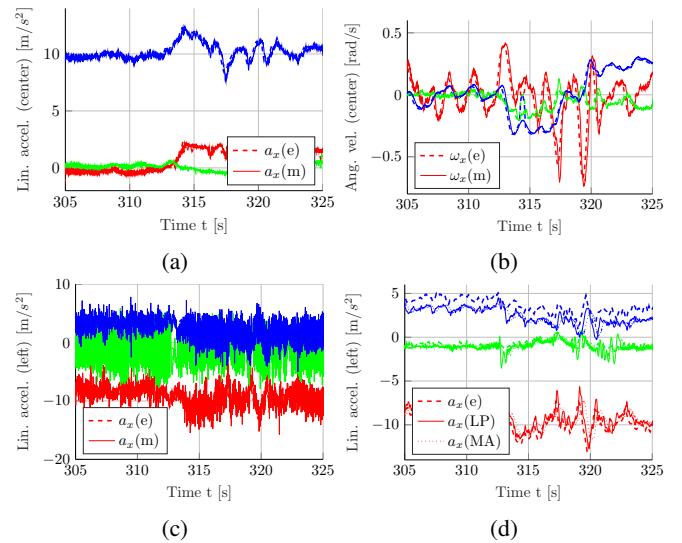


Fig. 6: Expected (e), measured (m) and filtered (low-pass: LP, moving average: MA) inertial measurements, recorded by the center and left wing tip IMU. The color coding is x (red), y (green), z (blue) throughout the paper.

not de-biased) inertial measurements for the center fuselage camera are shown in Fig. 6a, 6b. The IMU *ADIS16448* is used with a measurement sensitivity of $\pm 500^\circ/\text{s}$, FIR filter, and four filter taps. Fig. 6c shows the measured (raw, not de-biased) accelerometer readings for the left wing IMU using *MPU6050*'s on-board digital low-pass filter (DLPF) with a cut-off frequency $f_c = 94 \text{ Hz}$. Fig. 6d plots the solution from a low-pass filter with $f_c = 3 \text{ Hz}$ and moving-average filter with window size of 50 on the right. Based on these results, we conclude to use *MPU6050*'s on-board DLPF at $f_c = 5 \text{ Hz}$ (minimal available cut-off frequency) and compensate for the introduced constant, known delays.

B. Probabilistic wing model

Fig. 7 plots the measured relative transformation between the April tag attached to the right wing camera and the

right side camera $\mathbf{T}_{C_{\text{side},\text{right}}}^{\text{tag}_{\text{right}}}$. The experiment begins on the ground with take-off approximately at the 162.5 s mark. The relative translation and orientation vary within 5 cm and 6° respectively during the duration of the experiment (including on-ground detections). The take-off is particularly well observable in the rotation around the z -axis. The observed

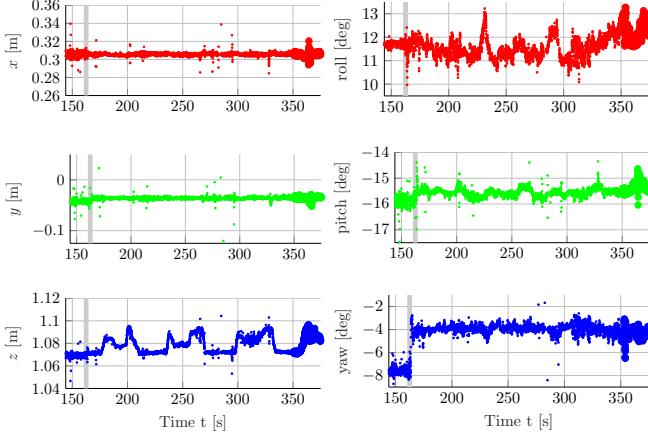


Fig. 7: Observations of the April Tag attached to the right wing camera, as seen from the right side camera. The grey vertical line marks the take-off (landing is not shown).

transformation $\mathbf{T}_{C_{\text{side},j}}^{\text{tag}_j}$ is then transformed to $\mathbf{T}_{B_c}^{B_j}$ via the rigid transformation chain given in Eq. 5 which is stored in the EKF in form of a mean and variance for the left-center and center-right stereo pair.

C. Vision update

Fig. 8 visualizes the quality difference between the relative camera-to-camera transformation priors, as well as the photometric alignment of the left to the center camera. By *Calib-ground* we denote the transformation estimate $\mathbf{T}_{C_c}^{C_l}$ obtained using *Kalibr* in a static setup and with no load applied to the wings. In contrast, *Calib-air* represents the mean of our wing model (cf. Sec. IV-C). Features tracked in the center camera are projected into the left camera according to Eq. 2: The features re-projected with the rigid relative transformation $\mathbf{T}_{C_c}^{C_l}$ obtained from *Calib-ground* are shown in green, the features re-projected using the mean of the wing model (*Calib-air*) are shown in red. One can

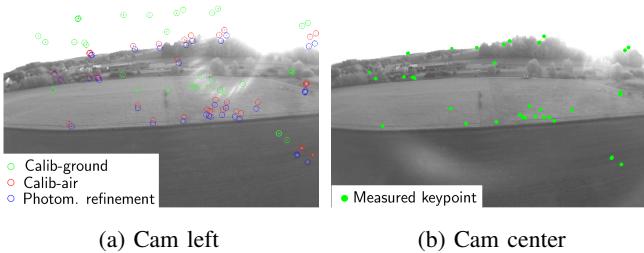


Fig. 8: Wing model prior and photometric refinement step for the alignment of the wing and center camera image.

observe that the features obtained from *Calib-ground* are far off from the correct feature position, while *Calib-air* returns a relative good initial position in particular during level flight. The feature positions refined by the photometric sparse image alignment are marked in blue (Eq. 3). The photometric formulation shows a robust performance even under challenging conditions such as the lense flare.

D. Depth map

Fig. 9 depicts the depth maps for the left-center stereo pair, computed using efficient block-matching (BM) [29], with $\mathbf{T}_{C_c}^{C_l}$ from b) *Calib-ground*, c) *Calib-air*, and d) our proposed framework. Areas with low texture variation can potentially be filled in by using different stereo pairs and by observing the area from different viewpoints.

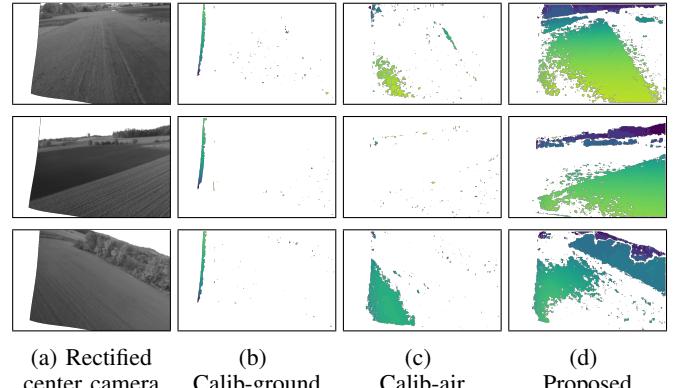


Fig. 9: Single shot depth maps obtained from the left and center image, expressed in the rectified center camera.

E. Runtime

The runtime results of the main modules are given in Table I, measured on an *Intel i7-4800MQ* CPU at 2.70 GHz for comparison. The combination of relative IMU integration,

Module	mean [ms]	Remark
Relative IMU integration (0-order)	0.03	Frame to frame
Photometric refinement	3.2	Per frame
Rectification, block-matching	38.1	Per stereo pair

TABLE I: Runtime results

good relative pose prior for the sparse photometric image alignment, and the employed EKF formulation makes the framework extremely efficient. With the tested frame rate of 10 Hz the image stream can be processed with some margin. Furthermore, the modular formulation makes it possible to run several visual-inertial stereo pairs in different threads. At this point, the employed rectification and block matching algorithm appears to become the first bottleneck if the image rate is to be set to a higher value.

VII. CONCLUSIONS AND FUTURE WORK

In this paper, we further developed the idea of using visual-inertial sensor systems on-board a UAV for improved

environmental awareness and demonstrated the effectiveness in real-world experiments. We investigated the challenges encountered in the three modules of the framework: The problem of measuring strong high-frequency vibrations on the wing tip IMUs was solved in software in form of a low-pass filter. In future work, we aim at a solution that minimizes the vibrations already on the hardware level. Since the relative transformation is estimated, a flexible mount or damping material could be used to isolate the visual-inertial system from the motor and wind gust induced vibrations encountered on the wing.

Furthermore, we described our procedure to obtain a probabilistic wing model, formulated as a Gaussian prior, and showed the discrepancy to the on-ground calibration. In future work, a more sophisticated wing model could be identified. For instance, the wing could be modeled as a cantilever beam with additional inputs such as air speed. As the deformation of the wing influences the lift distribution our model and accelerometer readings could also be incorporated into the on-board controller, e.g. for wind gust rejection.

Finally, in contrast to [1], we used a photometric sparse image alignment formulation to compute the relative transformation between wing and center camera. This approach has shown to be a robust way to transfer and track features from one to another camera under challenging conditions. In a next step, we intend to test our framework on a platform with stronger wing flapping behavior such as (manned) gliders.

ACKNOWLEDGMENT

We wish to thank: T. Stastny (experiments and discussion), T. Steiger and Y. Allenspach (safety pilots), A. Ruckli (robustification of April tag detection), T. Mantel and A. Melzer (help with VI-Sensor).

REFERENCES

- [1] T. Hinzmann, T. Taubner, and R. Siegwart, “Flexible stereo: Constrained, non-rigid, wide-baseline stereo vision for fixed-wing aerial platforms,” in *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 2550–2557, May 2018.
- [2] M. Bloesch, S. Omari, M. Hutter, and R. Siegwart, “Robust visual inertial odometry using a direct EKF-based approach,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 298–304, IEEE, 2015.
- [3] S. Leutenegger, S. Lynen, M. Bosse, R. Siegwart, and P. Furgale, “Keyframe-based visual-inertial odometry using nonlinear optimization,” *The International Journal of Robotics Research*, vol. 34, no. 3, pp. 314–334, 2015.
- [4] R. Mur-Artal and J. D. Tardós, “ORB-SLAM2: an open-source SLAM system for monocular, stereo and RGB-D cameras,” *IEEE Transactions on Robotics*, vol. 33, no. 5, pp. 1255–1262, 2017.
- [5] S. Houben, J. Quenzel, N. Krombach, and S. Behnke, “Efficient multi-camera visual-inertial SLAM for micro aerial vehicles,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 1616–1622, IEEE, 2016.
- [6] K. Sun, K. Mohta, B. Pfommer, M. Watterson, S. Liu, Y. Mulgaonkar, C. J. Taylor, and V. Kumar, “Robust stereo visual inertial odometry for fast autonomous flight,” *IEEE Robotics and Automation Letters*, vol. 3, no. 2, pp. 965–972, 2018.
- [7] P. Gohl, D. Honegger, S. Omari, M. Achtelik, M. Pollefeys, and R. Siegwart, “Omnidirectional visual obstacle detection using embedded FPGA,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 3938–3943, IEEE.
- [8] M. Müller, F. Steidle, M. J. Schuster, P. Lutz, M. Maier, S. Stoneman, T. Tomic, and W. Stürzl, “Robust visual-inertial state estimation with multiple odometries and efficient mapping on an MAV with ultra-wide FOV stereo vision,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 3701–3708, IEEE, 2018.
- [9] A. J. Barry and R. Tedrake, “Pushbroom stereo for high-speed navigation in cluttered environments,” in *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 3046–3052, IEEE, 2015.
- [10] M. Warren, D. McKinnon, and B. Upcroft, “Online calibration of stereo rigs for long-term autonomy,” in *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 3692–3698, IEEE, 2013.
- [11] M. Warren, P. Corke, and B. Upcroft, “Long-range stereo visual odometry for extended altitude flight of unmanned aerial vehicles.” *The Intern. J. of Robotics Research*, vol. 35, no. 4, pp. 381–403, 2016.
- [12] P. Lanier, N. Short, K. Kochersberger, and L. Abbott, “Modal-based camera correction for large pitch stereo imaging,” in *Structural Dynamics, Volume 3*, pp. 1225–1238, Springer, 2011.
- [13] P. Lanier, “Stereovision Correction Using Modal Analysis,” Master’s thesis, Virginia Polytechnic Institute and State University, 2010.
- [14] N. J. Short, “3-D Point Cloud Generation from Rigid and Flexible Stereo Vision Systems,” Master’s thesis, Virginia Tech, 2009.
- [15] P. Yang, X. Chen, and J. Wang, “Decoupling of airborne dynamic bending deformation angle and its application in the high-accuracy transfer alignment process,” vol. 19, p. 214, Multidisciplinary Digital Publishing Institute, 2019.
- [16] M. W. Achtelik, S. Weiss, M. Chli, F. Dellaert, and R. Siegwart, “Collaborative stereo,” in *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on*, pp. 2242–2248, IEEE, 2011.
- [17] L. Teixeira, F. Maffra, M. Moos, and M. Chli, “VI-RPE: visual-inertial relative pose estimation for aerial vehicles,” *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 2770–2777, 2018.
- [18] M. Karrer, M. Agarwal, M. Kamel, R. Siegwart, and M. Chli, “Collaborative 6DoF Relative Pose Estimation for two UAVs with Overlapping Fields of View,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2018.
- [19] M. Achtelik, S. Weiss, M. Chli, F. Dellaert, and R. Siegwart, “Collaborative Stereo,” in *Proceedings of the IEEE/RSJ Conference on Intelligent Robots and Systems*, 2011.
- [20] C. Forster, Z. Zhang, M. Gassner, M. Werlberger, and D. Scaramuzza, “SVO: semidirect visual odometry for monocular and multicamera systems,” *IEEE Trans. Robotics*, vol. 33, no. 2, pp. 249–265, 2017.
- [21] T. Hinzmann, T. Schneider, M. Dymczyk, A. Schaffner, S. Lynen, R. Siegwart, and I. Gilitschenski, “Monocular visual-inertial SLAM for fixed-wing UAVs using sliding window based nonlinear optimization,” in *Intern. Symp. on Visual Computing*, pp. 569–581, Springer, 2016.
- [22] T. Hinzmann, T. Schneider, M. Dymczyk, A. Melzer, T. Mantel, R. Siegwart, and I. Gilitschenski, “Robust map generation for fixed-wing UAVs with low-cost highly-oblique monocular cameras,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 3261–3268, IEEE, 2016.
- [23] C. Forster, M. Pizzoli, and D. Scaramuzza, “Svo: Fast semi-direct monocular visual odometry,” in *2014 IEEE international conference on robotics and automation (ICRA)*, pp. 15–22, IEEE, 2014.
- [24] A. Fusielo, E. Trucco, and A. Verri, “A compact algorithm for rectification of stereo pairs,” *Machine Vision and Applications*, vol. 12, pp. 16–22, Jul 2000.
- [25] J. Rehder, J. Nikolic, T. Schneider, T. Hinzmann, and R. Siegwart, “Extending kalibr: Calibrating the extrinsics of multiple IMUs and of individual axes,” in *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4304–4311, May 2016.
- [26] F. Furrer, M. Fehr, T. Novkovic, H. Sommer, I. Gilitschenski, and R. Siegwart, *Evaluation of Combined Time-Offset Estimation and Hand-Eye Calibration on Robotic Datasets*. Cham: Springer International Publishing, 2017.
- [27] 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 *IEEE International Conf. on Robotics and Automation (ICRA)*, pp. 431–437, 2014.
- [28] F. Sittel, J. Müller, and W. Burgard, “Computing velocities and accelerations from a pose time sequence in three-dimensional space,” Tech. Rep. 272, University of Freiburg, Department of Computer Science, 2013.
- [29] G. Bradski, “The OpenCV Library,” *Dr. Dobb’s Journal: Software Tools for the Professional Programmer*, vol. 25, no. 11, pp. 120–123, 2000.