Toward GPS-Denied, Multi-Vehicle, Fixed-Wing Cooperative Localization

Gary Ellingson* Tim McLain
Department of Mechanical Engineering
Brigham Young University
Provo, Utah

***Abstract*—This paper describes a vision and proposes a method for multiple, small, fixed-wing aircraft cooperatively localizing in GPS-denied environments. Recent work has focused on the development of a monocular, visual-inertial odometry for fixed-wing aircraft that accounts for fixed-wing flight characteristics and sensing requirements. The odometry was developed to be a front-end for novel methodology called relative navigation, which has been developed in prior work. This paper describes how the front-end could enable a back-end where odometry from multiple vehicles and inter-vehicle measurements could be used in a single, global, back-end, graph-based optimization. The inter-vehicle measurements over constrain the graph and allow the optimization to remove accumulated drift for more accurate estimates. The goal of this work is to show that many, small, potentially-lower-cost vehicles could collaboratively localize better than a single, more-accurate, higher-cost GPS-denied system.**

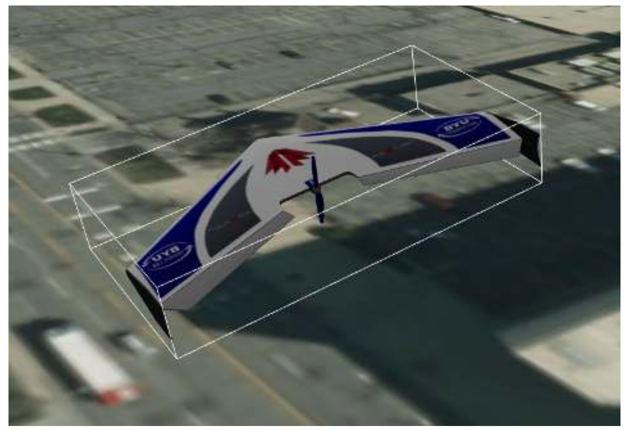

Fig. 1. This work enables GPS-denied navigation on fixed-wing aircraft. This high-fidelity, aircraft simulation was used to test the front-end, odometry estimator.

I. Introduction

More than ever before unmanned aircraft systems (UAS) need the ability to accurately navigate in GPS-denied environments. In both civil and defense applications UAS need to have an accurate knowledge of their motion to complete their mission objectives. The advent of highly-accurate, miniaturized navigation systems that fuse inertial measurements with GPS measurements (GPS-INS) have allowed UAS to operate in a variety of new applications. These navigation capabilities remain limited because GPS-INS solutions are brittle to GPS signal degradation and dropout. For example, civil autonomous drone delivery services will need to accurately navigate in and around obstacles where GPS signals are partially or fully obstructed.

Many military defense applications require aerial navigation in areas where GPS signals have been spoofed or jammed. Some applications require long-distance, high-speed flights and limited communication with ground-based command centers. In contrast to low-flying delivery and inspection aircraft, these vehicles require less precision because of their distance from obstacles, but need to limit the accumulation of drift over time to achieve their mission objective.

*gary.ellingson@byu.edu

UAS benefit from being small and inexpensive. Aircraft designers often speak of size, weight, and power (SWaP) constraints that influence trade-offs in the design. Navigation capabilities have similar constraints. GPS-denied solutions that only use inertial measurement units (IMUs) have been successfully implemented, but these solutions are only possible with highly-accurate, prohibitively-expensive, military-grade IMUs that have been precisely calibrated. Small UAS often must utilize sensors that are much less precise and instead use advanced algorithms to account for sensor noise and remove drift from state estimates. Constructing small, lower-quality vehicles make it possible to economically produce more vehicles to perform the mission rather than one, higher-quality vehicle.

In general small UAS can also benefit from collaboration to produce synergistic effects. Specifically for GPS-denied navigation, collaboration may provide significant advantages. Without position measurements to limit drift, global position and yaw angle are unobservable [1]–[3]. Vehicles must use other exteroreceptive sensing to help limit how fast estimate drift accumulates. If multiple vehicles could share measurements then the drift of all the vehicles could be further limited and provide even

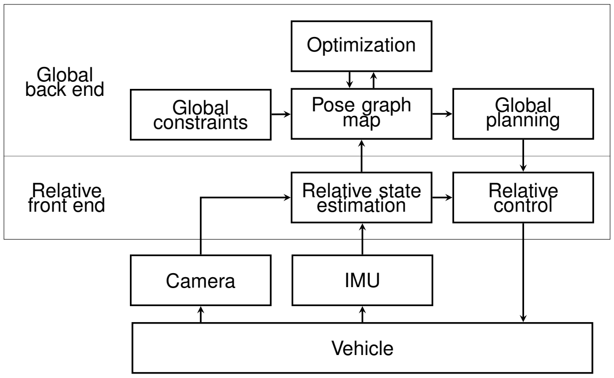

Fig. 2. The relative navigation architecture was developed for GPS-denied navigation. Estimation and control are performed in a front end where the vehicle operated relative to a local coordinate frame. The back end accounts for global information by utilizing odometry from the front end and optimizing it in a global pose-graph map.

better accuracy.

GPS-denied navigation on fixed-wing UAS requires specific sensing and estimation considerations. The majority of previous GPS-denied research and development has mainly focused on multirotor aircraft. Fixed-wing UAS differ from multirotors because they have different aircraft dynamics, they generally fly at higher speeds, and they are unable to stop and hover in place. Multirotor UAS are often able to utilize depth sensors, such as laser scanners, to effectively measure their motion because they can fly in and around structure in the environment. On fixed-wing UAS, depth sensors are less effective at measuring the motion of the aircraft because they usually fly high above the environment. This work proposes the development of a method to enable multiple, small, fixed-wing UAS to collaboratively localize.

II. Previous Works

This proposed work draws from previous research in three areas: The overall GPS-denied architecture utilizes the relative navigation framework, the front-end, visual-inertial odometry is a modification to the multi-state-constraint Kalman filter (MSCKF), and the back-end optimization comes from the wealth of literature on pose-graph optimization. Relevant work in these areas is summarized in the following sections.

A. Relative Navigation

Researchers have recently proposed on a new approach to GPS-denied navigation called relative navigation [4], [5]. It is a methodology and framework that separates the navigation into two sub-tasks. It separates a front-end estimator from a back-end optimization. The front end operates relative to the local surroundings and a back end that uses regular updates from the front end to create and maintain a global map. Figure 2 shows the framework architecture.

Relative navigation is motivated by a fairly simple concept called the relative-reset step [6] which is closely related to keyframe-based methods. The concept is for the front-end estimator to regularly declare a new local origin at the location of the vehicle. This also serves to remove uncertainty from the filter because the new origin is defined to be exact. At each new origin the prior transform can be sent to the back end as an edge in a directed pose graph.

The relative navigation approach has several advantages over contemporary methods. It is locally observable by construction and it has better filter consistency compared to other state-of-the-art approaches [7]. The front end has the computational advantages of an extended Kalman filter (EKF). The pose graph used in the back end is able to better represent large, nonlinear errors in odometry estimates. The back end can also incorporate other constraints, such as opportunistic GPS measurements or place-recognition loop closures.

Several tests have been perform to demonstrate relative navigation [8]. Assumptions about vehicle dynamics, sensing, and filtering have mostly limited the tests to multirotor aircraft at relatively low speeds. The approach has also be implemented with the entire architecture on a single vehicle that has enough computational resources. Sensing requirements have ensured the paths are in and around structured environments which have allowed the paths be relatively short and include loops back on themselves. These factors have limited the impact of the relative navigation architecture as a solution to the GPS-denied navigation problem.

B. Relative MSCKF

In the majority of the relative navigation work the front-end state estimator has been called the relative multiplicative extended Kalman filter (RMEKF) [6]. The RMEKF has required a keyframe-based odometry as a measurement and the odometries have used depth sensors, such as laser scanners and RGBD cameras, to resolve scale ambiguity. Fixed-wing aircraft, where RGBD and laser depth sensors are impractical due to the increased distance to features in the environment, require a different approach. Further, the main functions of the RMEKF were to combine inertial and visual odometry measurements and to perform a relative reset at each keyframe declaration. The odometry alone would otherwise be sufficient to provide the back end with odometry edge transformations from pose to pose.

More recently, a new tightly-couple, visual-inertial odometry has been introduced as a front-end estimator [9]. It uses only monocular imagery, without depth measurements, for exteroreceptive sensing. It combines the odometry calculations, inertial measurements, and relative-reset steps into one filter. This filter was developed specifically to enable fixed-wing UAS to use the relative navigation framework.

The new filter is based on the MSCKF. The MSCKF is more ideal for fixed-wing UAS because it makes no assumptions about the distance to observed features, requires no depth measurements, and makes no assumptions about the vehicle dynamics. The MSCKF uses a unique measurement model that was originally presented in [10]. It avoids adding uncertainty to the filter by not initializing states that are not well known. Further, updates are performed after a image feature moves out of view and all information about that feature is obtained.

Since its introduction, the MSCKF has seen extensive development in the literature. It has been demonstrated for use on ground vehicles [11], spacecraft [12], and even smart phones [13]. It has also been compared to several more-recent visual-inertial odometries and its accuracy and consistency properties remain comparable to the state-of-the-art with less computational burden [14].

Since the new, front-end filter was first presented in [9], several improvements have been introduced. The main improvements come from the reimplementation of the filter in the C++ programing language instead of Python. The greater speed from C++ allows the filter to run in real time while using more tracked features and more images per second. It was tested using the ROS/Gazebo simulation tools that were developed as part of ROSplane [15]. Figure 1 shows an example of the simulation. The filter now produces nearly double the accuracy as it did in previous results [9]. Figure 3 shows the trajectory of the aircraft and the accumulated estimates from the filter. Progress has also been made to use the algorithm in a hardware flight demonstration and results will be published when they are available.

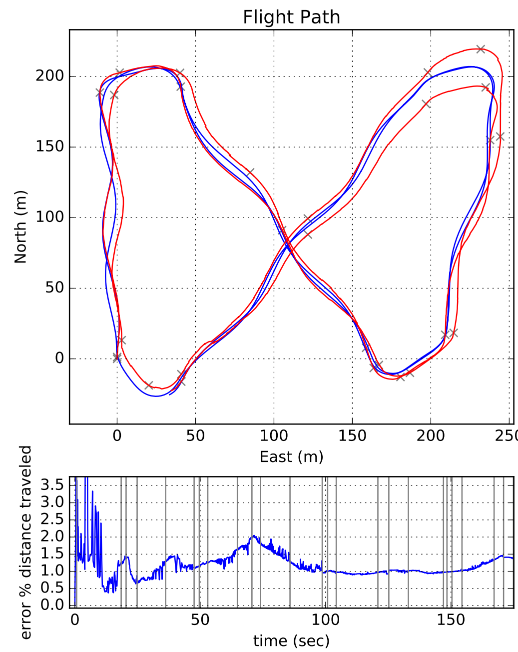

Fig. 3. Top: The path of the aircraft. The accumulated estimate (red) is compared actual path (blue). Gray × indicate relative-reset origins. Bottom: Accumulated error is less than 2% of the total distance traveled. Gray vertical lines indicated relative resets. The aircraft flew nearly 2200 m and the filter is nearly twice as accurate as previously reported results.

C. Graph Optimization

The relative-navigation back end has, in the past, been used to keep track of the global map by creating a directed pose graph. During the relative reset, the position and heading angle states and covariances are zeroed and the transformation from just before the reset is sent to the global back-end as an edge in the graph. Covariance uncertainty is effectively removed from the front-end filter and sent to a global back end where the pose graph has the ability to represent non linear uncertainties from yaw better than a Gaussian filter [7]. The back end is able to do edge optimization on the graph of pose estimates to improve global states for performing a global mission. The optimization is also able to incorporate and account for other constraints, such as opportunistic GPS measurements and place-recognition loop closures, for more accurate localization.

Graph-based optimization methods have been effectively used in robotic localization for some time [16]. Advances in computational power and sparse-matrix mathematics have, more recently, increased both the speed at which the optimizations can be performed and the number of nodes, or factors, that can be considered in the graph. Generalized graph optimization (g2o) [17], Georgia Tech smoothing and mapping (GTSAM), and incremental smoothing and mapping (ISAM) [18] are all graph optimization frameworks that have open-source implementations that are available for research. In the past, the relative-navigation back end has used the g2o graph optimization framework but recently other methods have been explored.

GTSAM is a smoothing and mapping toolbox that uses factor graphs to iteratively optimize a bipartite graph [19]. This means that it performs maximum a-posteriori inference through the relationships of states and factors that relate the states. Factors can be sensor measurements or odometry between aircraft poses. Odometry estimates are binary factors and measurements, such as opportunistic GPS or bearing to static features, are unary factors.

Because the global back end uses a pose graph that is a relatively sparse representation of the vehicle odometry, it has potential to be useful for multi-vehicle cooperative localization. Multiple vehicle cooperation has the potential to limit estimate drift over extended flights due to the increased baseline between sensors [20], [21]. Other work has show that multiple vehicles can collaboratively estimate using poses as factors in factor graph smoothing frameworks [22].

GTSAM can be easily applied to new problems. It allows implementing factors for new measurement models by inheriting from a factor class and implementing the loss

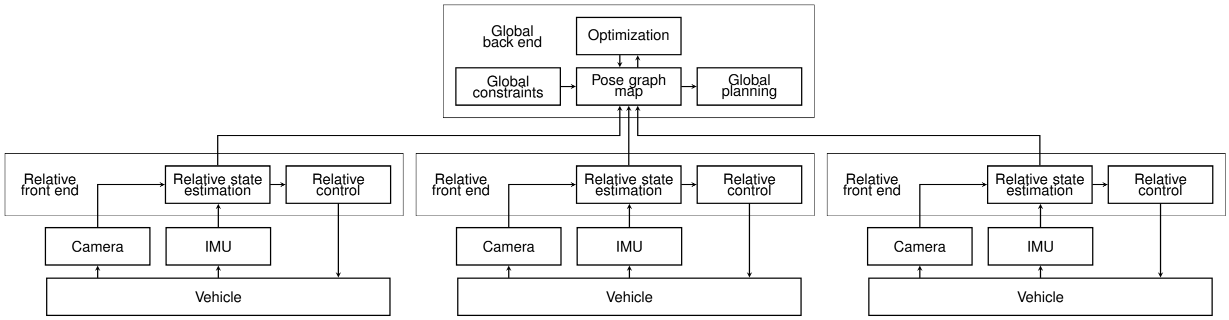

Fig. 4. The proposed, multi-vehicle architecture will include each vehicle with a separate front-end odometry that provides edges to centralized back end. This work will initially focus on localization in the back end and not on global path-planning or relative control.

function for the new factor [19]. This functionality will be necessary for using GTSAM as a multi-vehicle back end with custom inter-vehicle measurements.

III. Development

We propose the creation of a new relative-navigation back-end graph and optimization that will incorporate the odometry and measurements from multiple vehicles. The back end will use the new visual-intertial odometry that was constructed for use of fixed-wing aircraft. The aircraft will be able to fly high above the environment over relatively long, straight flights. The vehicles will also have the ability to measure the distance to the other UAS through inter-vehicle range measurements. These measurements, with odometry, will be combined in a centralized back-end graph. Figure 4 shows the modified relative-navigation architecture with a centralized, global back end. Figure 5 shows the odometry transformations and inter-vehicle measurements that make up the graph of the UAS flight paths.

This work will use GTSAM framework for optimizing the back-end pose graph. Relative transformations from the front-end filter will be used as binary factors between consecutive aircraft poses. Multiple-aircraft localization will be accomplished by incorporating pose variables and odometry factors from multiple aircraft in the same optimization. The distance measurements between aircraft can be modeled by implementing a binary factor with the distance between aircraft poses as the factor constraints. Initially, this work will assume distance measurement are only taken simultaneous to the front-end, relative reset but then will continue by relaxing that assumption through the addition of nodes in the graph that correspond to the timing of the measurement. A similar approach that was used in [5] to utilize multiple GPS measurements between nodes.

An intermediate step to the full multi-vehicle problem will be to incorporate measurements to known landmarks. These landmarks may be visual landmarks or stationary transceivers that provide distance measurements. In this case the measurements are unary factors in the graph and the entire graph must shift in the global reference frame rather than being relative to the starting pose of the aircraft. This step will help in the development of the back-end while still providing a valuable contribution since these type of measurements are feasible in many flight scenarios.

One potential challenge that this work may encounter is the difficulty of optimizing vehicle poses with large initial uncertainties. The odometry of the aircraft will tend to accumulate position error over time and the error may become large for longer-distance flights. When large uncertainties exist, pose-graph optimizations can get stuck in local minima that are far from the optimal solution. Recent work in the BYU MAGICC Lab has shown robustness to initialization errors can be achieved by optimizing the edge transformations rather then the poses of the vehicle. This robustness comes, however, at some increased computational cost. The work in [23] also directly deals with large initial uncertainty for optimizing graphs with inter-vehicle range measurements.

IV. Experimental Setup

The method will initially be demonstrated in simulation where the full system can be simulated in detail. Figure 6 shows a Gazebo simulation with three independent aircraft flying in close proximity. The simulation will be invaluable for development and testing because experiments can be performed with relative ease and truth comparison is both possible and simple. The inter-vehicle range measurements can be implemented as sensor simulations plugins.

The architecture will use ROS for communication and messaging. The use of ROS will make the move from simulation to hardware flight tests easier because simulated sensors can be replaced with hardware sensors on the UAS.

The flight-test experiments will take place on small fixed-wing platforms. A suitable aircraft platform is the STRIX StratoSurfer by Ready Made RC. This aircraft works well because it is sturdy and has a large payload capability. Some initial testing has already been done on this aircraft. The aircraft will carry an Odroid single-board computer and use ROSplane [15] for the aircraft

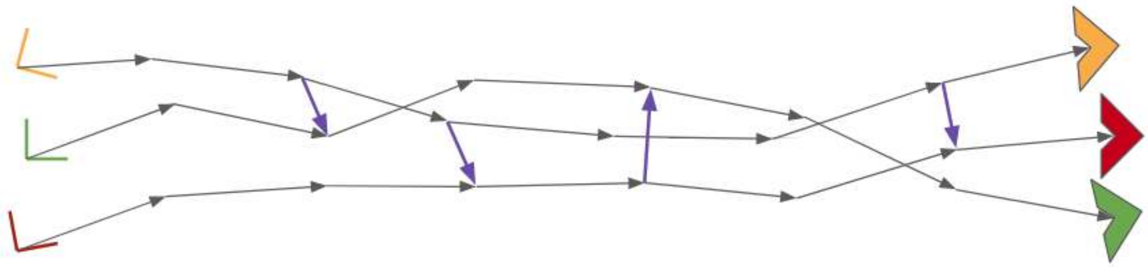

Fig. 5. A pose graph will be optimized in the global back end. It will demonstrate relative-navigation back end can be used for multi-vehicle collaborative localization. The demo will have multiple fixed-wing aircraft fly over relatively long, straight flights where loop-closure like measurements are obtained from inter-vehicle range measurements. Odometry is represented by black arrows and inter-vehicle measurement are purple arrows.

Fig. 6. The proposed method will be demonstrated in a high-fidelity simulation where multiple aircraft can be simulated, each with independent front-end estimators and inter-vehicle measurements.

stabilization and autonomous control. The true aircraft odometry will be measured by an accurate GPS-INS implementation for comparison.

V. Conclusion

This work will demonstrated a feasible method for collaboratively localizing fixed-wing UAS in GPS-denied environments. The work will be significant because it will directly acknowledge and address challenges of GPS-denied, fixed-wing UAS. GPS-denied solutions for multi-rotor aircraft are fairly common, but less so for fixed-wing aircraft. Often, when solutions do exist, the approaches make significant simplifying assumptions, such as operating over flat-earth or in Manhattan world environments, or having complex or unreasonable sensing requirements, such as downward facing camera or depth measurements. The completion of this work uses minimal sensing (only camera, IMU, and inter-vehicle range) and makes no such simplifying assumptions. It further enables GPS-denied navigation within a collaborative framework capable of incorporating inter-vehicle measurements from multiple aircraft. These measurement will over constrain the graph and allow the graph smoothing and optimization to remove accumulated error from the graph.

The work will also extend the impact of the relative navigation framework. It will allow the value of relative navigation to be shown for a different type of vehicle with a different mission profile. Since the back end constrains the graph with inter-vehicle measurements and not with loop closures, the aircraft will be able to fly in relatively long, straight paths at high speeds. These mission profiles may be more representative of real-world UAS scenarios.

The approach will be demonstrated using a Gazebo simulation of a small fixed-wing aircraft with simulated sensors. The simulated aircraft dynamics and sensor-noise characteristics will be representative of those from an actual small, unmanned fixed-wing aircraft. The accuracy and consistency of the relative odometry are presented, as well as hardware results with GPS as ground truth comparison.

This research will contribute to the maturation of small unmanned aircraft. Before introduction into the national airspace or use in military applications, small unmanned aircraft will need greater reliability and to be robust to GPS signal degradation and dropout. This research will utilize state-of-the-art methods and modify and combine them in novel ways to expand the capabilities of these aircraft.

This work will include simulation and hardware flight testing to demonstrate the capabilities of the proposed methods. The multi-aircraft cooperative flight demonstration that is enabled by the new relative-navigation back end and the tightly-coupled, visual-inertial front end, will show the value of the complete system with relative navigation. It will be successful if multiple aircraft can cooperatively localize with greater accurately than an individual aircraft.

Acknowledgment

Gary Ellingson receives generous funding from the Utah NASA Space Grant Consortium Fellowship.

This work has been funded by the Center for Unmanned Aircraft Systems (C-UAS), a National Science Foundation Industry/University Cooperative Research Center (I/UCRC) under NSF award No. IIP-1161036 along with significant contributions from C-UAS industry members.

REFERENCES

[1] S. Weiss, M. W. Achtelik, S. Lynen, M. Chli, and R. Siegwart, "Real-time onboard visual-inertial state estimation and self-calibration of MAVs in unknown environments," in *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pp. 957–964, IEEE, 2012.

[2] E. Jones, A. Vedaldi, and S. Soatto, "Inertial structure from motion with autocalibration," in *Workshop on Dynamical Vision*, vol. 25, 2007.

[3] T. D. Barfoot and P. T. Furgale, "Associating uncertainty with three-dimensional poses for use in estimation problems," *IEEE Transactions on Robotics*, vol. 30, no. 3, pp. 679–693, 2014.

[4] R. C. Leishman, T. W. McLain, and R. W. Beard, "Relative navigation approach for vision-based aerial GPS-denied navigation," *Journal of Intelligent and Robotic Systems*, vol. 74, no. 1–2, pp. 97–111, 2014.

[5] D. O. Wheeler, P. W. Nyholm, D. P. Koch, G. J. Ellingson, T. W. McLain, and R. W. Beard, "Relative navigation in GPS-degraded environments," *Encyclopedia of Aerospace Engineering*, 2016.

[6] D. P. Koch, D. O. Wheeler, R. Beard, T. McLain, and K. M. Brink, "Relative multiplicative extended Kalman filter for observable GPS-denied navigation," 2017.

[7] D. Wheeler, D. Koch, J. Jackson, T. McLain, and R. Beard, "Relative navigation: An observable approach to GPS-degraded navigation," *IEEE Control Systems, Accepted*, 2017.

[8] D. O. Wheeler, D. P. Koch, J. S. Jackson, G. J. Ellingson, P. W. Nyholm, T. W. McLain, and R. W. Beard, "Relative navigation of autonomous GPS-degraded micro air vehicles," 2017.

[9] G. Ellingson, K. Brink, and T. McLain, "Relative visual-inertial odometry for fixed-wing aircraft in GPS-denied environments," in *Position, Location and Navigation Symposium (PLANS), 2018 IEEE/ION*, IEEE, 2018.

[10] A. I. Mourikis and S. I. Roumeliotis, "A multi-state constraint kalman filter for vision-aided inertial navigation," in *Robotics and automation, 2007 IEEE international conference on*, pp. 3565–3572, IEEE, 2007.

[11] L. E. Clement, V. Peretroukhin, J. Lambert, and J. Kelly, "The battle for filter supremacy: a comparative study of the multi-state constraint Kalman filter and the sliding window filter," in *Computer and Robot Vision (CRV), 2015 12th Conference on*, pp. 23–30, IEEE, 2015.

[12] A. I. Mourikis, N. Trawny, S. I. Roumeliotis, A. E. Johnson, A. Ansar, and L. Matthies, "Vision-aided inertial navigation for spacecraft entry, descent, and landing," *IEEE Transactions on Robotics*, vol. 25, no. 2, pp. 264–280, 2009.

[13] M. Shelley, *Monocular visual inertial odometry on a mobile device*. PhD thesis, Masters thesis, Technischen Universitat Munchen, 2014.

[14] J. Delmerico and D. Scaramuzza, "A benchmark comparison of monocular visual-inertial odometry algorithms for flying robots," *Memory*, vol. 10, p. 20, 2018.

[15] G. Ellingson and T. McLain, "ROSplane: Fixed-wing autopilot for education and research," in *Unmanned Aircraft Systems (ICUAS), 2017 International Conference on*, 2017.

[16] S. Thrun, W. Burgard, and D. Fox, *Probabilistic robotics*. MIT press, 2005.

[17] R. Kummerle, G. Grisetti, H. Strasdat, K. Konolige, and W. Burgard, "g2o: A general framework for graph optimization," pp. 3607–3613, May 2011.

[18] B. Kim, M. Kaess, L. Fletcher, J. Leonard, A. Bachrach, N. Roy, and S. Teller, "Multiple relative pose graphs for robust cooperative mapping," in *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, pp. 3185–3192, IEEE, 2010.

[19] F. Dellaert, "Factor graphs and GTSAM: A hands-on introduction," tech. rep., Georgia Institute of Technology, 2012.

[20] R. Sharma, S. Quebe, R. W. Beard, and C. N. Taylor, "Bearing-only cooperative localization," *Journal of Intelligent & Robotic Systems*, vol. 72, no. 3-4, p. 429, 2013.

[21] R. Sharma and C. Taylor, "Cooperative navigation of MAVs in GPS denied areas," in *Multisensor Fusion and Integration for Intelligent Systems, 2008. MFI 2008. IEEE International Conference on*, pp. 481–486, IEEE, 2008.

[22] L. R. Sahawneh and K. M. Brink, "Factor graphs-based multi-robot cooperative localization: A study of shared information influence on optimization accuracy and consistency," 2017.

[23] A. Charkraborty, K. Brink, R. Sharma, and L. Sahawneh, "Relative pose estimation using range-only measurements with large initial uncertainty," *under-review*, 2017.