2024 IEEE International Conference on Robotics and Automation (ICRA 2024)
May 13-17, 2024. Yokohama, Japan

Measurement-limited Multi-Agent, Relative Pose Estimation for On-Orbit Inspection

Mark Mercier, *Member, IEEE,* David Curtis, Clark Taylor, *Senior Member, IEEE*

***Abstract*— Relative navigation methods are a critical enabling technology for the next generation of autonomous spacecraft conducting close proximity operations. This is especially true for multi-agent inspection operations in which safety including intra-agent or agent-target collisions are a serious concern. Additionally, in an on-orbit servicing operation various failure modes of the target may result in unreliable *a-priori* knowledge or cooperation from the target. The main contribution of this work is the demonstration of a method for multi-agent, relative pose estimation that is robust to A) sensor blinding and B) dynamic uncertainty. This objective is accomplished leveraging GTSAM, an existing toolbox for the formulation of factor graphs, along with an algorithm for the efficient, real-time solution of such factor graphs, iSAM2. This estimation method is demonstrated in an example scenario with uncertain dynamics and sensor blinding due to sun position. Results revealed that the iSAM2-based method is capable of handling sensor blinding through leveraging an inter-agent range measurement, despite a dynamically uncertain environment.**

I. Introduction

As a part of the larger guidance, navigation, and control (GNC) architecture of a spacecraft, relative navigation methods are crucial in enabling spacecraft operations. This holds especially true in spacecraft rendezvous and proximity operations (RPO)—like in-space servicing and debris removal—that often leverage autonomy to compensate for ground communication delays or dropouts. In both cases, inspection to gain knowledge of the target is required as a prerequisite for action.

Most literature concerned with relative pose estimation in space do so either via simultaneous localization and mapping (SLAM) techniques, by using *a-priori* knowledge, or with fiducials [1], [2], [3], [4]. This work assumes no *a-priori* knowledge of the RSO is available and hardware on board is not capable of reliable, real-time image processing or resolved imagery—imagery in which the target is taking up significant pixels so the shape can be resolved—is unavailable. Removing dependence on resolved imagery enables the methods demonstrated herein to be applied with less restriction on distance to target or sensor quality.

Multiple works use a distributed, filter-based [5], [6] or centralized, factor graph (FG) [7] approach to perform cooperative localization (CL)—translation-only estimation—in a multi-agent, RPO scenario. Additionally, [8] performs multi-agent, angles-only CL with consideration for sensor drop-outs; however, it does not leverage an inter-agent range constraint and focuses more on dynamics due to the much longer duration of operations. A smoothing-based approach rather than a pure filtering approach has been shown to yield improved results in various scenarios with degraded measurements [9], [10]. Two recent works utilize a smoothing-based approach to conduct 6-DoF pose estimation of an unknown target in single and multi-agent RPO scenarios [11], [12]. The use of resolved imagery, a range to the target, and inertial measurements in both works separate them from the contributions of this work.

During any multi-agent mission, the capability of agents to sense, communicate, and fuse measurements determines the information available to a centralized navigation algorithm. The centralized method in this work uses knowledge from all agents simultaneously, rather than a distributed approach that operates with only information available on the edge—a single inspecting agent. The implementation of this centralization could be an on-orbit hub [13] or all agents to provide redundancy; however, the details of this centralization scheme are outside the scope of this work. Assuming a centralized estimation scheme provides a more complete picture of the inspection scenario through the incorporation of measurements from all sensors and is often considered to provide the best case estimate [14].

In an on-orbit inspection, certain sensors are required to accomplish the primary mission, such as an optical sensor for collecting images of a target, a data link for communication between agents, and internal attitude odometry systems—typically combined with the attitude control system. If only measurements derived from these sensors can be used for navigation, further size, weight, and power requirements on spacecraft design can be avoided. Although multi-agent close proximity operations are relatively unproven in the space domain, most theoretical and planned work points to autonomy as a requirement to conduct operations due to the generally slow flow of information between space assets and ground operators [15]. For this reason, a data link is expected to be a requirement between inspectors and timing information can provide an accurate and consistently available range measurement. Previous literature has demonstrated the impact of unreliable communication channels on multi-agent systems [16], but this work assumes constant, regular communication is available.

Additionally, a passive optical sensor is often the preferred choice for an inspection spacecraft due to its simplicity and low energy requirement. The use of such a sensor adds a requirement to consider the solar lighting condition, which has the potential to blind the optical sensor during the scenario. The use of purely relative measurements may

M. Mercier, D. Curtis, and C. Taylor are with the Air Force Institute of Technology, Wright-Patterson AFB, OH, 45420 USA e-mail: mark.mercier.4@spaceforce.mil.

enable application of this estimation method in dynamically uncertain environments—where inertial parameters such as the orbital parameters of the RSO are not precisely known. Dynamic model uncertainty has been discussed on a much smaller scale for application to on-orbit robotics[17] but not in the context of overcoming this uncertainty by leveraging uniquely multi-agent measurements.

The main contribution of this work is the development of a method for multi-agent, relative pose estimation using only measurements expected to be readily available from existing sensors without requiring resolved imagery or computationally expensive image processing techniques. The satisfaction of this objective leads to a method capable of enabling safe multi-agent, autonomous spacecraft RPO via accurate state estimation with minimal system requirements. In addition, special consideration is given in demonstrating robustness against the potentially troublesome on-orbit scenarios of sensor blinding and dynamic uncertainty.

The paper is outlined as follows. Section II provides background information on the relative dynamics model used in this work and a brief overview of FGs. Section III describes the methodology that establishes a simulation environment and the structure of the FG. Section IV provides results and analysis demonstrating the efficacy of the applied method and analysis on the impact of sensor blinding and dynamic uncertainty. Section V concludes the report and provides some avenues for future work.

II. Background

The background for this work begins with an abridged derivation of an existing relative spacecraft dynamics model that is leveraged in order to define a relative frame centered around the RSO. The section continues with a brief overview of FGs and standard solution methodologies, followed by a description of some of the features unique to iSAM2 as implemented through GTSAM [18].

A. Dynamics

The on-orbit scenario used in this work takes place in a geostationary orbit. By leveraging the circular nature of this orbit, a suitable dynamics model can be simplified from the general 2-body problem (2BP) equations of motion to the circular nonlinear equations of relative motion (cNERMs) so that the propagation of the state of a spacecraft can be measured with respect to a reference orbit. In this work, the term reference orbit is used to describe the RSO's inertial orbit around the Earth relative to which the inspector's orbits are determined. This simplification results in a constant mean motion ($\dot{\theta} = n$) where θ is the true anomaly of the RSO. The cNERMS can be seen in (1) and are used to generate the simulated truth trajectory, enabling measurement generation and results error analysis [19].

$$\begin{aligned} \ddot{x} - 2n\dot{y} - n^2 x &= -\frac{\mu(r_c + x)}{||\vec{r}_d||^3} + \frac{\mu}{r_c^2} + a_x \\ \ddot{y} + 2n\dot{x} - n^2 y &= -\frac{\mu y}{||\vec{r}_d||^3} + a_y \\ \ddot{z} &= -\frac{\mu z}{||\vec{r}_d||^3} + a_z \end{aligned} \tag{1}$$

In (1) x, y, and z refer to the position of an inspector along the local-vertical, local-horizontal (LVLH) axes, $||\vec{r}_d|| = \sqrt{(r_c + x)^2 + y^2 + z^2}$ is the distance of the inspector from the origin of the Earth-centered inertial (ECI) frame, r_c is the radius of the circular reference orbit, μ is the gravitational parameter, and a_x, a_y, and a_z are external disturbances, such as thrust or environmental perturbations. The LVLH frame is inconsistently defined throughout literature, but in this case with the RSO in a circular orbit, the origin is assumed to be co-located with the RSO and $\hat{x}$, $\hat{y}$, and $\hat{z}$ are aligned with the inertial orbit's radial, velocity, and normal to the orbit plane directions respectively as shown in Fig. 1 relative to the ECI frame. Note that a critical assumption is made that the attitude of the RSO is arbitrary and the pose of each agent is relative to the LVLH frame whose origin is fixed at the RSO, but whose axes are not necessarily aligned with those of the RSO body axes.

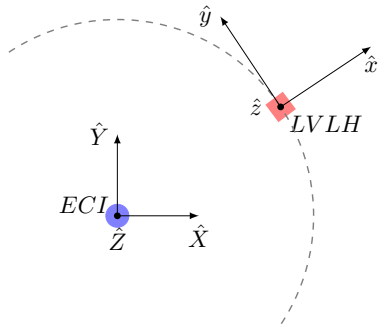

Fig. 1. Visualization of LVLH-frame as defined for cNERM dynamics

To ensure the estimation methods developed in this work are robust to any real-world disturbances, random perturbations are included during simulated trajectory generation. Although these perturbations do not accurately simulate the perturbations expected on-orbit in terms of bias, they do reflect the magnitude of disturbances encountered on-orbit—which is sufficient for testing estimator robustness in short time duration scenarios. Notably, no dynamic factor is included in the FG structure (detailed further in Section III.B) and subsequent states are only connected via attitude odometry. This choice means that although the dynamic model used to propagate the truth trajectory for this scenario is assuming a circular orbit, nothing inherent in the estimation method limits it to only this application. A demonstration of the robustness of this method to dynamically uncertain situations is a focus of the results to follow.

B. Estimation Method

A FG is a bi-partite graph consisting of nodes—variables to be estimated—and factors—prior information and measurements. This batch processing method attempts to find the *maximum a-posteriori* (MAP) estimate for a set of variables ($\boldsymbol{X}$). Optimal values ($\boldsymbol{X}^*$) are found for the estimated variables by maximizing the joint probability when also considering the noisy measurements. Through the common assumption that all measurements are Gaussian, an optimization problem in terms of error residuals can be written as in (2).

$$\boldsymbol{X}^* = \underset{\boldsymbol{X}}{\operatorname{argmin}} \sum_k ||\mathbf{z}_k - h(\boldsymbol{X}_k)||^2_{\Sigma_k} \tag{2}$$

where $\mathbf{z}_k$ is the measurement value, $h(\boldsymbol{X}_k)$ is the expected measurement corresponding to the state $\boldsymbol{X}_k$, $||\cdot||^2_\Sigma \equiv \cdot^T \Sigma^{-1} \cdot$ is the Mahalanobis distance, and Σ_k is the measurement covariance [11]. The MAP estimate based on the full FG solution is the optimal estimate and used throughout literature as the truth for scenarios where truth data is unavailable [20]; however, the solution time for a full FG grows unbounded in time with the number of included measurements [21].

This work focuses on iterative optimization methods, which are generally more computationally expensive than filtering methods, but are provably convergent [22] and often result in improved performance in highly nonlinear applications [20]. A Gauss-Newton method is a common choice for solving nonlinear, least squares problems that extends Newton's method to iteratively approximate the minimum of the sum shown in (2). A core component of this method is the need for an initial guess at which to calculate the initial residual and derivatives. Determining an appropriate initial guess can be challenging for problems where the trajectory of the states is not well understood [23] *a-priori* and may result in convergence to a local optimum [24]. A potential solution to this problem is iteratively expanding the time horizon of estimation and using previous iterations as initial guesses for the newly added time steps. These two initial guess techniques are compared in a dynamically unpredictable environment in the results to follow.

An efficient solver such as iSAM2 [25] can be used to mitigate the computational demands of an iterative approach [26] through the use of a Bayes tree. Advantages of the Bayes tree, when combined with sparse nonlinear incremental optimization, results in improvements in efficiency by eliminating the need for periodic batch steps. Unique to iSAM2, when compared to traditional FG methods of nonlinear optimization, is operation directly on the graphical model rather than first converting the model to a sparse matrix. By taking this approach, iSAM2 is able to incrementally solve nonlinear estimation problems, with the objective of achieving a real-time approach. The emphasis on a real-time approach differs from a full FG MAP estimate, which often focuses on achieving the best estimate of a trajectory after the conclusion of a scenario. A full description of the incremental inference and elimination approach used in iSAM2 is outside the scope of this paper, but can be found in the original work [25].

III. Methodology

Unfortunately, real-world data sets for on-orbit inspections are not widely available and often lab-based experimentation suffers from inaccuracies due to difficulties in simulating relevant parts of the space environment (i.e. lighting conditions, distances, timing, and dynamics). For this reason, a simulation environment is often used that is able to sufficiently mimic on-orbit dynamics, while incorporating constraints such as sensor blinding due to sun lighting constraints. This section begins with a description of the simulation environment developed for this work and is followed by an outline of the structure of the FG, which includes some description of the real-world instrumentation expected to provide each measurement factor.

A. Simulation Environment

The true trajectories for three agents as well as derived measurements from this trajectory are generated by a simulation environment via the perturbed, cNERM dynamic model described in Section II.A. Dynamic model uncertainty is simulated via random perturbations and an adjustment of the mean motion of the reference orbit, essentially simulating a scenario in which the radius of the target's orbit differs from expected. The initial conditions for each agent are based on a common baseline trajectory for inspection in literature, natural motion circumnavigation (NMC)—a kind of passive, relative orbit that theoretically enables an agent to follow a 2:1 ellipse around a target. Some simplifying assumptions on reference orbit properties are embedded in this theory based on the Hill-Clohessy-Wiltshire equations of motion; however, these are outside the scope of this work [27]. Injection into different NMCs—a specific initial condition that results in passively following the ellipse around the target—is chosen as the initial state for each agent and some noise is added based on expected imperfect spacecraft thrusters on-orbit.

Measurements derived from this trajectory are simulated and imported by time step—mirroring on-orbit collection operations—to an estimation script that updates the FG structure and optimizes estimated state values in an online fashion. Sensor blinding is incorporated by "flagging" the generated measurement to an out of bounds value to trigger the estimation script to drop the measurement. This measurement dropout is implemented by creating a phantom measurement from current state estimates. This method of handling dropout is required due to observability limitations and does affect covariance estimation, but can be compensated by boosting the covariance associated with these phantom measurements.

B. Factor Graph Structure

The nodes for this inspection scenario are chosen to represent the RSO's position—fixed at the origin of the relative frame via a strict prior factor for ease of implementation—and the 6-DoF relative pose of the three inspecting agents

throughout time. A prerequisite to a fully observable system is ensuring sufficient measurement information is available to define the 18 states—6 states per agent for three agents—at each time step. Note that having enough measurements is necessary for observability, but is not sufficient and an observability guarantee is related to the linear independence of the information provided by each measurement.

In this work, we seek to have a set of factors composed only of unresolved, relative measurements using instrumentation expected to be available on-board an inspecting spacecraft. Four contributing sources of information meet these requirements, and sufficiently constrain the system nodes. First, prior information on the state of each agent as well as the assumption that the target remains at the origin of the relative frame is encoded in prior factors (P_{Ax}).

Second, each agent is expected to be equipped with a passive, optical sensor that provides a camera factor by determining the pixel location of the centroid of the target—or in the case of a distant target, a single pixel representing the entire target—in the sensor focal plane. This results in 6 measurements—2-D pixel location (ϕ_x, ψ_x) in the body frame of each of the three agents—that help to relate the position of each agent to the RSO and thus the relative frame. At each update step, the camera factor measurement is checked for sensor blinding in which case it is negated.

Third, a range measurement is assumed to be available between agents, but not between agents and the RSO—providing three additional measurements (D_{xy}). Fourth, instrumentation for attitude odometry, such as a gyroscope—providing three axis orientation data for three agents yielding 9 measurements (O_x)—is relatively common onboard spacecraft and is expected to be standard on an inspection spacecraft. Four measurement types—priors, camera, inter-agent range, and angular odometry—meet the previously described necessary condition for observability; however not redundantly and so are susceptible to degraded measurement quality in the event of measurement drop out. The robustness of this FG structure and the iSAM2 algorithm to a loss of some of these measurements is analyzed in the results to follow. A schematic showing the first two time steps of the FG can be seen in Fig. 2.

IV. Results

The results section begins with a description of the problem setup, including relevant parameters such as measurement covariances and system constraints. The section continues with some analysis on handling sensor blinding by leveraging inter-agent range. The section concludes with some discussion on dynamic uncertainty and comparing errors for a baseline estimate based on theory and expected initial conditions, two full FG estimates based on different initialization criteria, and a real-time iSAM2 estimate.

A. Problem Setup

Most parameters in Table I are self-explanatory; however, the parameters related to dynamic model uncertainty are worth further explanation. Recall that the cNERM dynamic model used in this work to propagate the state of each agent does not naturally account for any on-orbit perturbations, so random perturbations are added to simulate potential trajectory deviations. The dynamic perturbation covariance defines the zero mean normal distribution from which the orbital perturbations are sampled and is roughly the same magnitude as expected perturbations from solar radiation pressure, the prevailing orbital perturbation in geostationary orbit. The RSO Orbit Difference is the semi-major axis difference between the true inertial reference orbit—used to generate the true trajectory—and the expected inertial reference orbit—used to generate the baseline estimate that serves as the initialization for *a-posteriori* estimation. In general, results are sensitive to these parameters and scenarios with increased noise levels occasionally did not result in a converged estimate. This sensitivity illustrates the need to consider system capabilities when evaluating this estimation method, especially in scenarios where sensor blinding may occur.

TABLE I
Scenario Parameters

Name	Value (Units)
Prior Position Covariance	10^2 (m)
Prior Orientation Covariance	10^{-1} $(^\circ)$
Range Covariance (σ_D)	0.01^2 (m)
Camera Covariance	10.0^2 (pixels)
Gyro Covariance	0.005 $(\frac{^\circ}{sec})$
Dynamic Perturbation Covariance	10^{-7} $(\frac{m}{s^2})$
RSO Orbit Difference	2000 (km)
Time Between Measurements	60 (s)

The initial estimate based on the assumed problem dynamics is referred to as the baseline estimate and is the best estimate possible for the trajectory of each agent without considering measurements obtained during operations. This baseline serves as the initial guess for the *a-posteriori* full FG estimation and is dynamically imperfect for two reasons. First, the expected initial condition of each agent differs from actual due to an imperfect NMC insertion. Second, the trajectory is propagated via the unperturbed, cNERM dynamic model, which differs due to added random perturbations and uncertain reference orbit parameters.

Two primary error metrics are considered in evaluating the efficacy of the estimator. First, errors in position estimation are simply calculated as the Euclidean distance between the estimated position and the associated true position at the same time step. Second, the quality of the estimate of agent orientation is evaluated based on differences in the sensor boresight vector. As an example of why this is an important metric, consider an inspecting satellite that has a narrow FOV camera and the estimated sensor boresight differs significantly from the true sensor boresight. In this case, the attitude control on-board the spacecraft may force the attitude to change to a point where the inspector no longer has visibility of the target due to the poor estimate.

In addition to a baseline estimate derived from expected

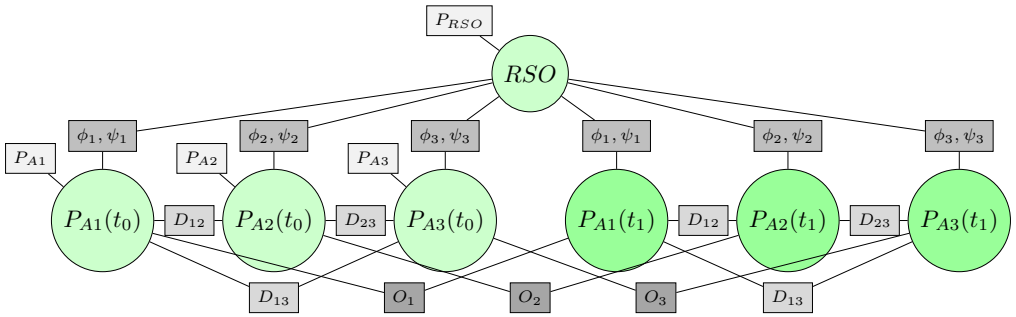

Fig. 2. Schematic showing the first two time steps of FG including nodes (green circles) and factors (grey squares)

initial conditions and dynamics, a full FG, MAP estimate is found using GTSAM. This full estimate theoretically represents the ideal estimate of each agent trajectory. Although this method represents the optimal estimate, it may be susceptible to converging to a local minimum in scenarios where a quality initial guess is unavailable. In contrast, iSAM2-based estimation is intended to run in real-time with estimates updating as measurements are collected—as would happen in operations.

B. Results Analysis

Results analysis can begin by demonstrating that the iSAM2-based algorithm can handle sensor blinding and dynamic uncertainty—as dictated by the parameters in Table I. In Fig. 3, the optimal case of "No Sensor Blinding"—where camera measurements are assumed to be constantly available—is shown to be unaffected by the sensor blinding zones as expected. In all plots, the colored, shaded region represents the time during which the correspondingly colored agent—as can be seen in the legends of Figs. 4 and 5—is sensor blind. The first method of handling sensor blinding is leveraging a persistent, inter-agent range measurement unique to a multi-agent operation. The "Degraded Range" method shows performance of the estimator for a scenario where the inter-agent range measurement is degraded by raising both the real and assumed covariance ($\sigma_D = 10m$). By incorporating a quality inter-agent range measurement, the estimator as described in this work is more robust to sensor drop out, but still suffers from slightly degraded performance during camera measurement unavailability.

The consistency of the estimator—its ability to predict its own error—is considered in Table II, which states the average of the three agents' Root Mean Squared Error (RMSE) and Average Normalized Estimation Error Squared (ANEES) for the position error of each method corresponding to Fig. 3. Note the ANEES for the first two methods is less than the translational degrees of freedom (3) indicating an estimated covariance larger than the actual error. In contrast, the degraded range measurement case has a larger RMSE and ANEES showing a general breakdown in the estimation method when a quality range measurement is unavailable.

Having demonstrated an iSAM2-based, real-time capable estimator robust to sensor blinding by leveraging uniquely multi-agent measurements, the second objective of ensuring robustness in the presence of dynamic uncertainty can be considered. First, the difficulty in handling an uncertain dynamic model can be seen as the difference between "FullFG" and "FullFG-Iterative" in the position error plot shown in Fig. 4. Recall that both methods use a Gauss-Newton method that depends on an initial guess to evaluate initial residuals and derivatives; however, they differ in their initial guess. "FullFG" is initialized using a guess based on an assumed dynamics model, shown in Fig. 4 as "BL". This poor initial guess causes the "FullFG" solution to converge to a local optimum as the scenario grows in length and the initial guess diverges from the true state. In contrast, an online method typically uses either the estimate from the previous time step or that estimate propagated forward via a dynamic model. The prior approach is taken by "FullFG-Iterative", which iterates through the scenario in a pseudo-online fashion using previous estimates as the initial estimate for newly added

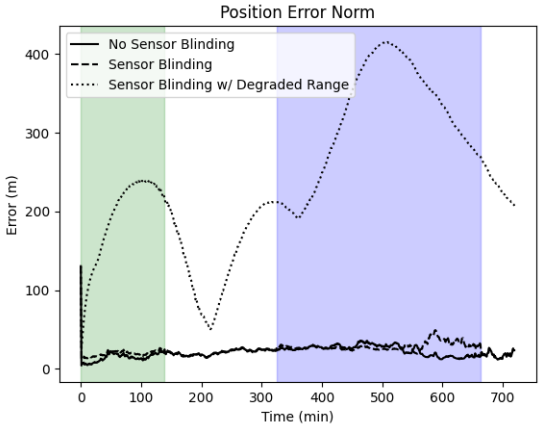

Fig. 3. Average agent position error between estimate and true trajectory from iSAM2 algorithm comparing handling of sensor blinding scenario

TABLE II
Average Position Error Metrics

Method	**RMSE**	**ANEES**
No Sensor Blinding	23.019	1.802
Sensor Blinding	26.240	1.973
Sensor Blinding w/ Degraded Range	277.312	191.688

nodes. This initial estimate technique appears to lead to a stable estimate even in the presence of an uncertain dynamic model as the initial guess is updated throughout the scenario. Importantly, this iterative initial estimate method may suffer in scenarios where time steps are significantly separated due to the implicit assumption of little motion from one time step to the next. This was not explored in this work but is unlikely to be an issue at normal measurement acquisition rates due to the slow relative dynamics of on-orbit RPO.

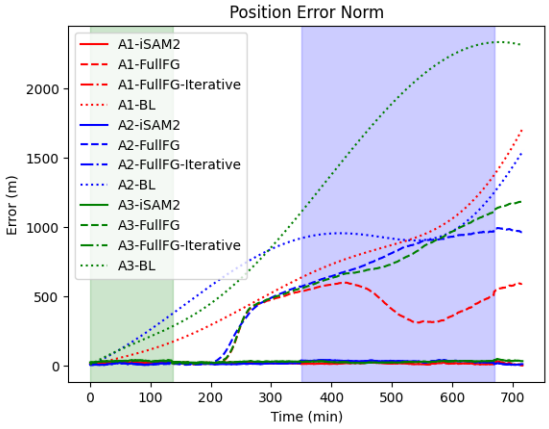

Fig. 4. Euclidean distance between estimated and true trajectory for baseline, full FG, and iSAM2 measurement

Recall that the expected available systems for inspection spacecraft do not provide measurements connecting the translational state of an agent from one time step to the next—traditionally done via translational odometry, IMUs, or dynamic factors. This deficiency leads to a difficulty in determining an initial guess—as well as generally weak observability—and while the iterative initial estimate approach used in "FullFG-Iterative" does help to remove dependence on an assumed dynamic model, it also causes increased computation time. Specifically, for estimating over the entire scenario duration as seen in Fig. 4, the computation time grows from 1.0s to 254.2s between the "FullFG" and "FullFG-Iterative". This drastic increase in computation time once again reveals that although the *a-posteriori* estimate represents the theoretically optimal solution, it comes at an exponentially increasing computation cost as the FG grows—in this case from 0.003s on the first iteration to 0.39s on the last. Fig. 5 plots the sensor boresight error and shows a poor estimate of the attitude in the "FullFG" method due to the coupling of the attitude estimate to the poor translational estimate due to the camera factor.

Fortunately, leveraging an online method like iSAM2, which is built to add measurements and progressively optimize as a scenario evolves, addresses both the initial guess and computation time difficulties. In Figs. 4 and 5, the iSAM2 solution and the "FullFG-Iterative" solution are nearly identical in terms of error metrics, but the iSAM2 version is able to achieve this at computation times much closer to those of the "FullFG" solution.

In summary, the iSAM2-based estimator described in this

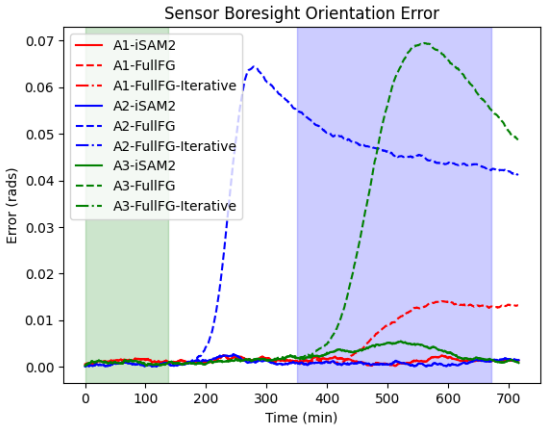

Fig. 5. Angular error between estimate and true sensor boresight for full FG and iSAM2 estimation methods

work accomplished the main contributions of being robust to A) sensor blinding and B) dynamic uncertainty. Robustness to sensor blinding is handled by providing a persistent inter-agent range measurement that achieves results similar to the ideal scenario of no sensor blinding. In addition, an online method based on leveraging an iterative initial estimate with a full FG showed robustness to dynamic uncertainty, albeit at the cost of computation time.

V. Conclusion

The objective of this work was to demonstrate an estimator capable of conducting relative pose navigation for multi-agent RPO using only relative, unresolved measurements. In addition, The unique demands of the space environment require an estimation method intended for application to autonomous systems to be robust to sensor blinding and dynamic uncertainty. This was accomplished by incorporating measurements derived from on-board instrumentation expected to be available on an inspecting spacecraft. A centralized approach was taken in which measurements from all spacecraft were available to simultaneously estimate all agent states. A FG was built and solved using both a full FG, *a-posteriori* approach and iSAM2. Results revealed the success of the iSAM2-based estimator in sensor blinding scenarios through the use of quality inter-agent range measurements. Difficulties in generating the initial estimate for an *a-posteriori* method were discussed and overcome through the use of a pseudo-online method that improved the estimate in a dynamically uncertain environment, but at the cost of increased computation time. In contrast, the online, iSAM2 estimator was shown to be capable of handling a dynamically uncertain environment as initialization was not dependent on any assumed scenario dynamics. The use of the GTSAM toolbox—and knowledge of its heritage on small processing capability-limited processors—facilitates future work, which may include Processor-In-the-Loop testing or lab experimentation in a sufficiently space-like environment.

REFERENCES

[1] D. Zou, P. Tan, and W. Yu, *Collaborative visual SLAM for multiple agents:A brief survey*. KeAi Communications Co., 10 2019, vol. 1, no. 5.

[2] B. E. Tweddle, "Relative computer vision based navigation for small inspection spacecraft," in *AIAA Guidance, Navigation, and Control Conference 2011*. American Institute of Aeronautics and Astronautics Inc., 2011.

[3] D. Fourie, B. E. Tweddle, S. Ulrich, and A. Saenz-Otero, "Flight results of vision-based navigation for autonomous spacecraft inspection of unknown objects," *Journal of Spacecraft and Rockets*, vol. 51, no. 6, pp. 2016–2026, 11 2014.

[4] M. Sabatini, R. Volpe, and G. Palmerini, "Centralized visual based navigation and control of a swarm of satellites for on-orbit servicing," *Acta Astronautica*, vol. 171, pp. 323–334, 2020.

[5] X. Xu, Q. Yuan, B. Song, P. Zheng, and Y. Xiao, "Cooperative localization algorithm of multiple spacecraft with angle-only measurements," in *Proceedings of 2021 5th Chinese Conference on Swarm Intelligence and Cooperative Control*. Springer, 2022, pp. 764–772.

[6] H. K. Sipowa and J. McMahon, "Distributed estimator for spacecraft cooperative localization," *Journal of Guidance, Control, and Dynamics*, vol. 45, pp. 1166–1171, 2022.

[7] M. Mercier and D. Curtis, "Relative navigation methods for a multi-agent, on-orbit inspection mission," *2023 IEEE/ION Position, Location and Navigation Symposium, PLANS 2023*, pp. 1303–1310, 2023.

[8] J. Sullivan, A. W. Koenig, J. Kruger, and S. D'Amico, "Generalized angles-only navigation architecture for autonomous distributed space systems," *Journal of Guidance, Control, and Dynamics*, vol. 44, pp. 1087–1105, 6 2021.

[9] I. Mohammed, S. J. Geetha, S. S. Shinde, K. Rajawat, and S. Chakrabarti, "Modified re-iterated kalman filter for handling delayed and lost measurements in power system state estimation," *IEEE Sensors Journal*, vol. 20, pp. 3946–3955, 4 2020.

[10] M. J. García-Ligero, A. Hermoso-Carazo, and J. Linares-Pérez, "Estimation from a multisensor environment for systems with multiple packet dropouts and correlated measurement noises," *Applied Mathematical Modelling*, vol. 45, pp. 802–812, 5 2017.

[11] T. P. Setterfield, "On-Orbit Inspection of a Rotating Object Using a Moving Observer," *Dissertation, Massachusetts Institute of Technology*, 2017.

[12] A. Terán Espinoza, W. D. Sanchez, and D. Roascio, "Sensor fusion for multi-agent spacecraft proximity operations via factor graphs," in *2018 AIAA SPACE and Astronautics Forum and Exposition*. AIAA, 2018.

[13] S. Silvestrini, A. Capannolo, M. Piccinin, M. Lavagna, and J. G. Fernandez, "Centralized autonomous relative navigation of multiple spacecraft around small bodies," *AIAA Scitech 2020 Forum*, vol. 1.F, 2020.

[14] J. M. Maley and R. Zurakowski, "Largest ellipsoid applications in decentralized multi-agent state estimation," in *AIAA SCITECH 2022 Forum*, 2022, p. 0754.

[15] S. Mian, T. Garrett, A. Glandon, C. Manderino, S. Balachandran, C. A. Munoz, and C. V. Dolph, "Autonomous spacecraft inspection with free-flying drones," *AIAA/IEEE Digital Avionics Systems Conference - Proceedings*, vol. October, 10 2020.

[16] Y. G. Sun and L. Wang, "Consensus of multi-agent systems in directed networks with nonuniform time-varying delays," *IEEE Transactions on Automatic Control*, vol. 54, pp. 1607–1613, 2009.

[17] S. Abiko and G. Hirzinger, "Adaptive control for a torque controlled free-floating space robot with kinematic and dynamic model uncertainty," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2009, pp. 2359–2364.

[18] F. Dellaert, "Factor graphs and gtsam: A hands-on introduction," *Georgia Institute of Technology, Tech. Rep*, vol. 2, p. 4, 2012.

[19] K. T. Alfriend, S. R. Vadali, P. Gurfil, J. P. How, and L. S. Breger, *Spacecraft Formation Flying*. Elsevier Astrodynamics Series, 2010.

[20] T.-C. Dong-Si and A. I. Mourikis, "Motion tracking with fixed-lag smoothing: Algorithm and consistency analysis," in *2011 IEEE International Conference on Robotics and Automation*. IEEE, 2011, pp. 5655–5662.

[21] H. Johannsson, M. Kaess, M. Fallon, and J. J. Leonard, "Temporally scalable visual slam using a reduced pose graph," in *2013 IEEE International Conference on Robotics and Automation*. IEEE, 2013, pp. 54–61.

[22] G. Sibley, G. S. Sukhatme, and L. H. Matthies, "The iterated sigma point kalman filter with applications to long range stereo." in *Robotics: Science and Systems*, vol. 8, no. 1, 2006, pp. 235–244.

[23] H. Fang, N. Tian, Y. Wang, M. Zhou, and M. A. Haile, "Nonlinear bayesian estimation: From kalman filtering to a broader horizon," *IEEE/CAA Journal of Automatica Sinica*, vol. 5, pp. 401–417, 3 2018.

[24] G. C. Goodwin, J. C. Aguero, and R. E. Skelton, "Conditions for local convergence of maximum likelihood estimation for armax models," *IFAC*, pp. 771–776, 2003.

[25] M. Kaess, H. Johannsson, R. Roberts, V. Ila, J. Leonard, and F. Dellaert, "iSAM2: Incremental Smoothing and Mapping Using the Bayes Tree," Tech. Rep., 2011.

[26] J. Hsiung, M. Hsiao, E. Westman, R. Valencia, and M. Kaess, "Information sparsification in visual-inertial odometry," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 1146–1153.

[27] E. R. Prince, "Optimal Finite Thrust Guidance Methods for Constrained Satellite Proximity Operations Inspection Maneuvers," *Dissertation, Air Force Institute of Technology*, 2018.