



Generalized Angles-Only Navigation Architecture for Autonomous Distributed Space Systems

Joshua Sullivan,*[✉] Adam W. Koenig,[†] Justin Kruger,[‡] and Simone D'Amico[§]
Space Rendezvous Laboratory, Stanford, California 94305

<https://doi.org/10.2514/1.G005439>

This paper addresses the development and verification of an estimation architecture for autonomous relative navigation in multi-satellite systems using angles-only measurements. Unlike traditional angles-only navigation approaches, which are generally hindered by poor handling of observability constraints, this prototype solution is designed to be flexible and applicable to a multitude of relevant missions. The estimator converges accurately and robustly in varying orbital environments without requiring maneuvers by exploiting system nonlinearities using an unscented Kalman filter that is streamlined by formulating the measurement and dynamics models using relative orbital elements. While the majority of prior work has focused on only solving for the relative motion of a single target, the approach here enables estimation of several additional parameters. Key design trades are evaluated for estimating additive sensor biases, target ballistic properties in environments with strong nonconservative perturbations, adaptive process noise statistics, and the absolute orbit of the observing satellite. The architecture is generalized for relative navigation with multiple targets, and for decentralized estimation using multiple coordinated observers. These filter strategies are used to inform the development of a novel on-orbit demonstration of angles-only navigation known as the Starling Formation-flying Optical eXperiment (StarFOX) in partnership with NASA. Key hardware-in-the-loop algorithm verification results for StarFOX-specific scenarios are provided.

I. Introduction

FOLLOWING from the successes of recent flagship spacecraft formation-flying demonstrations, including GRACE (NASA) [1], TanDEM-X (DLR) [2], PRISMA [3], and the Magnetospheric Multiscale (MMS) mission (NASA) [4], active research and development efforts are now transitioning to spacecraft swarm concepts [5–10]. Swarms promise an expansion of the benefits offered by traditional low-scale formations, including additional redundancy from single-point failures, mission flexibility, and new domains of application. In particular, new missions for deep space exploration, on-orbit assembly of large structures, and distributed sensing using sparse aperture arrays or distributed antennae are all enabled by spacecraft swarms. These platforms, however, will be largely characterized by a duality of requiring improved situational awareness and navigation capability to avoid collision, while also consisting of smaller, cost-effective satellites with minimal computational and sensing capabilities. Accordingly, spacecraft of these envisioned formation-flying and swarm concepts must satisfy increasingly strict requirements on dynamics, guidance, navigation and control accuracy, autonomy in early phases of the mission, and resource efficiency. With these constraints in mind, this paper addresses the development of a generalized navigation architecture using a vision-based estimation technique known as angles-only navigation, culminating in the design of a novel on-orbit experiment that will enable future multi-satellite missions flying in Earth orbit and beyond.

In angles-only navigation, an observing spacecraft seeks to estimate the relative orbital state of target space objects that appear only as small clusters of pixels in the camera field of view. These pixel centroids intrinsically describe line-of-sight vectors that point from the observer's camera frame to the targets and are each subtended by a set of two angles, denoted as the bearing angles. In the context of autonomous satellite swarms, the use of vision-based sensing is motivated by the fact that it provides a robust, high dynamic-range, and passive navigation capability that uses simple sensors that are already on board most spacecraft. Furthermore, because of their low cost, low power consumption, and small form factor, vision-based sensors (VBSs) enable accurate relative navigation while also complementing the current trend of spacecraft miniaturization.

Despite these many advantages, navigation algorithms designed to use angles-only measurements incur several distinct challenges. First and foremost, limited dynamical observability resulting from using 2D bearing angles makes the complete estimation of the 6D relative orbital motion of a target space object difficult or impossible [11–13], particularly when using linear dynamics and measurement models. When angles-only estimation is successful, the convergence timescale is on the same order as the orbital period. Conducting known orbital maneuvers of the observing spacecraft has become a common solution employed in the literature to yield filter convergence [12,14–16], because doing so generates a known variation in the bearing angle trends that can be exploited to reconcile the unobservable relative range. However, this approach has the undesirable effect of strongly coupling the achievable relative navigation accuracy with the maneuver-planning task and must be repeated periodically to correct filter estimate divergence. Furthermore, the majority of research studies have been limited to low Earth orbit (LEO) applications, and generally involve scenarios between only two satellites. A noteworthy exception to the latter comes from the research of Wang et al. [17], wherein a consensus Kalman filtering strategy is pursued using measurements of a single target from multiple observers. Although the analytic backing for that approach is strongly conveyed, the simplified simulation testing against a minimalistic truth model and without extensive consideration of realistic noise and error sources leaves much to be proven about its practical performance capability in actual on-orbit applications. As a final note of deficiency in the current literature, improper treatment of the limited dynamical observability inherent to angles-only navigation has generally resulted in estimation architectures with limited applicability beyond basic orbital state estimation; additional useful parameters like sensor biases, ballistic properties,

Received 3 June 2020; revision received 31 October 2020; accepted for publication 12 December 2020; published online 23 April 2021. Copyright © 2021 by the authors. Published by the American Institute of Aeronautics and Astronautics, Inc., with permission. All requests for copying and permission to reprint should be submitted to CCC at www.copyright.com; employ the eISSN 1533-3884 to initiate your request. See also AIAA Rights and Permissions www.aiaa.org/randp.

*Stanford Aeronautics & Astronautics Department; also Senior Guidance, Navigation, and Control Engineer, Astranis Space Technologies, San Francisco, CA; jsullivan0926@gmail.com.

[†]Postdoctoral Researcher, Stanford Aeronautics & Astronautics Department, Durand Building, 496 Lomita Mall.

[‡]Doctoral Candidate, Stanford Aeronautics & Astronautics Department, Durand Building, 496 Lomita Mall.

[§]Assistant Professor, Space Rendezvous Laboratory Director, Stanford Aeronautics & Astronautics Department, Durand Building, 496 Lomita Mall.

observer absolute orbital state, and mismodeled (or completely unknown) maneuvers are simply not accounted for at all.

Angles-only navigation has been demonstrated in two major on-orbit experiments. The first of these, the Advanced Rendezvous Demonstration using Global Positioning System and Optical Navigation (ARGON) [18], was successfully conducted in 2012 during the PRISMA mission. Using ground-based image processing, relative orbit determination, and maneuver planning, ARGON demonstrated non-cooperative rendezvous from 30 km to as close as 3 km of separation between the two satellites. The next breakthrough angles-only experiment was the Autonomous Vision Approach Navigation and Target Identification (AVANTI) [19–21] phase of the FireBird mission, conducted in late 2016. Unlike ARGON, AVANTI conducted all image processing, state estimation, and maneuver planning on board and achieved rendezvous to within tens of meters of a non-cooperative ejected nanosatellite. While ARGON relied on prior knowledge of the orbital states from NORAD two-line elements (TLEs), AVANTI tested a new target identification and filter initialization algorithm, which reduced reliance on prior ground-based knowledge. Finally, AVANTI demonstrated navigation robustness while coping with substantially more data outages due to sensor blinding and thruster actuation constraints than anticipated. It is worthwhile to note that, just like the majority of research studies found in the literature, these two on-orbit experiments involved only binary formations in LEO and required extensive maneuver planning and actuation to effectively handle the observability constraints of angles-only navigation.

In response to the many challenges posed by angles-only navigation, Sullivan et al. [13] and Sullivan and D’Amico [22] developed a maneuver-free approach using the unscented Kalman filter (UKF) that was verified for applications in near-circular LEO and highly elliptical orbit (HEO). The method proposed in those works uses the relative orbital element (ROE) parameterization of the dynamics for several key strengths. First, the weak observability is decoupled into a single state element that well-approximates the range. Second, the filter state propagation is streamlined by using an efficient and accurate linear state transition matrix developed by Koenig et al. [23], which captures secular and long-period effects of the J_2 perturbation. And third, the dynamical observability is improved by exploiting nonlinearities relating mean to osculating ROEs (and then to bearing angles) in the UKF measurement model. These nonlinearities capture separation-dependent features of the relative motion that disambiguate the weakly observable range. The net result is a filtering methodology that accurately converges to an estimate of the full relative motion state without requiring orbital maneuvers. Still, that work left open several research avenues to be pursued, including the possibility for a more accurate estimation approach using numerical dynamics propagation, inclusion of simplified maneuver strategies, estimating additional useful parameters including the absolute orbit of the observer, applications involving multiple targets and/or observers, and more rigorous Monte Carlo testing.

This paper builds upon [13,22] and the legacy of the ARGON and AVANTI flight experiments by combining new breakthroughs in algorithm design and implementation with the development of a novel on-orbit formation-flying experiment using angles-only navigation. In particular, this work provides the following contributions to the state-of-the-art. First, several key improvements to the existing navigation framework are developed in order to enable new applications, and to provide greater performance and strategic design flexibility to future mission planners. In particular, the method of dynamics modeling is now more deeply treated as a design parameter wherein potential observability, accuracy, robustness, and efficiency can be traded off for the intended application by considering both linear and nonlinear state propagation in the filter. In a similar way, the success of the “maneuver-free” sequential filter design previously presented in [22] shows that orbital maneuvering is no longer a strict requirement for navigation convergence. Instead, simple single-impulse maneuvering schemes, as opposed to the complex observability-optimizing control algorithms often presented to enable angles-only navigation [14–16], can be used as a means of improving estimation convergence time and accuracy; maneuvering is now a strategic design parameter, rather than

a rigid requirement for navigation stability. As a means of extending the filter efficacy in myriad applications, this work rigorously assesses the impact of two complementary strategies, namely, process noise adaptation and state augmentation. In the former, the bearing angle measurement residuals are used to vary the filter tuning in real-time for improved performance during high-uncertainty or anomalous events. The latter enables the filter to estimate additional parameters that are useful to the particular scenario, such as sensor biases for a poorly calibrated camera, ballistic properties for an unknown target object, and the observer absolute orbit. State augmentation is also used to expand the maneuver-free angles-only estimator application from binary observer-target formations to scenarios with multiple targets and/or cooperative observers. Rigorous algorithm verification via Monte Carlo simulation allows filter design trades to be weighed against the requirements of the mission developer.

Second, a full systems engineering process is conducted to design a comprehensive angles-only navigation flight experiment, which reflects these new changes to the angles-only estimation paradigm. The result is the Starling Formation-flying Optical eXperiment (StarFOX), a navigation payload and series of tests designed to advance the technology readiness of swarm relative navigation for future deep space applications. The experiment will be hosted on the NASA Starling-1 formation [24] consisting of four satellites launched into LEO in 2022. StarFOX will provide valuable on-orbit experience for the first ever demonstration of adaptive angles-only nonlinear estimation involving multiple observer and target satellites in both maneuver-free and maneuvering scenarios.

Following this introduction, Sec. II is devoted to modeling preliminaries that focus on the relative orbital dynamics, the angles-only measurement model, and high-fidelity orbit propagation. Section III provides a detailed discussion of the angles-only navigation algorithms and simulation results for the sequential relative orbit estimation strategy. The application of these relative navigation strategies in an on-orbit demonstration is highlighted in Sec. IV with the introduction of StarFOX and specific hardware-in-the-loop simulation results. Finally, Sec. V offers lessons learned, conclusions, and ways forward with open research questions.

II. Modeling Preliminaries

As a precursor to the development of the angles-only navigation framework, a few key preliminaries must first be discussed. In particular, the following section will highlight two of the main ingredients for any state estimator: the dynamics and measurement models. This is then followed by a brief description of the high-fidelity orbit propagation utility used in this work and a summary of the simulation verification pipeline that will show up in several subsequent sections.

A. Relative Orbital Dynamics Models

This section provides preliminary fundamentals for modeling the relative motion dynamics between multiple space objects. Rather than providing an exhaustive discussion, the insight provided here is a high-level overview of key concepts that will be tailored to the application of angles-only relative navigation. For a comprehensive resource on spacecraft relative motion dynamics models, the work of Sullivan et al. [25] is recommended. In the context of relative orbital motion, the dynamics model parameterizes the trajectory of an orbiting space object with respect to a reference orbit. The reference orbit may describe the physical orbit of another space object or some other meaningful virtual orbit. For this work, the observer spacecraft orbit is treated as the reference about which the target space object motion is described. Irrespective of the dynamical state, the relative orbital mechanics are generally described by a system of nonlinear differential equations:

$$\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t), \mathbf{u}(t), \mathbf{p}(t), t) + \mathbf{w}(t) \quad (1)$$

where $\mathbf{x}(t) \in \mathbf{R}^n$ is the relative motion state, $\mathbf{u}(t) \in \mathbf{R}^m$ is the control input vector applied by either the observer or the target, $\mathbf{p}(t)$ is a vector of dynamical parameters (e.g., absolute orbit state, force-modeling parameters, etc.), and $\mathbf{w}(t) \in \mathbf{R}^n \sim \mathcal{N}(\mathbf{0}, \mathbf{Q})$ is the additive process

noise vector with symmetric positive semidefinite covariance matrix $Q \in \mathcal{S}_+^{n \times n}$. The nonlinear equations in Eq. (1) may be linearized about the reference state and control input to produce the linear dynamical state equations:

$$\dot{\mathbf{x}}(t) = A(t)\mathbf{x}(t) + B(t)\mathbf{u}(t) + \mathbf{w}(t) \quad (2)$$

In the above, $A(t) \in \mathcal{R}^{n \times n}$ denotes the time-variant plant matrix, which captures the unforced dynamics. Instead, the time-variant input sensitivity matrix, $B(t) \in \mathcal{R}^{n \times m}$, captures the forced dynamics due to $\mathbf{u}(t)$. In this linearized system, $A(t)$ and $B(t)$ are simply the Jacobian matrices of the vector-valued function \mathbf{f} with respect to $\mathbf{x}(t)$ and $\mathbf{u}(t)$, respectively. In general, the nonlinear dynamics model in Eq. (1) must be solved using numerical integration. Instead, it is often possible to find closed-form solutions to the model in Eq. (2) that propagate the state from time t_j to time t_k as

$$\mathbf{x}(t_k) = \Phi_{k,j}\mathbf{x}(t_j) + \int_{t_j}^{t_k} \Phi_{k,\tau}[B(\tau)\mathbf{u}(\tau) + \mathbf{w}(\tau)]d\tau \quad (3)$$

The zero-input response and zero-state response are accounted for in terms of the state transition matrix (STM), $\Phi \in \mathcal{R}^{n \times n}$, via propagation and convolution with the control input matrix, respectively. The choice of dynamics model [i.e., Eq. (1) or Eq. (2)] and solution methodology [numerical integration or linear propagation via Eq. (3)] are major design choices for the estimation architecture and are largely driven by the selection of the dynamical state representation. In the majority of research studies in the literature, the state is parameterized using either translational state components or using ROEs.

1. Translational State Elements

The relative translational state is made up of the relative position $\delta\mathbf{r}^{\mathcal{R}} = (x, y, z)$ and relative velocity $\delta\dot{\mathbf{r}}^{\mathcal{R}} = (\dot{x}, \dot{y}, \dot{z})$ of the target defined in the rotating reference frame \mathcal{R} centered on the observer. The \mathcal{R} frame is often denoted as the local-vertical/local-horizontal (LVLH) or radial/along-track/cross-track (RTN) frame. It is composed of the orthogonal basis vectors $\hat{\mathbf{R}}$ directed along the observer absolute position vector, $\hat{\mathbf{N}}$ in the direction of the observer orbital angular momentum vector, and $\hat{\mathbf{T}} = \hat{\mathbf{N}} \times \hat{\mathbf{R}}$, which completes the right-handed triad. The evolution of this relative state is given by a direct application of Newton's law of gravitation, resulting in the fundamental relative orbital differential equations [26]:

$$\begin{aligned} \delta\ddot{\mathbf{r}}^{\mathcal{R}} = & -\mu \left(\frac{\mathbf{r}^{\mathcal{R}} + \delta\mathbf{r}^{\mathcal{R}}}{\|\mathbf{r}^{\mathcal{R}} + \delta\mathbf{r}^{\mathcal{R}}\|_2^3} - \frac{\mathbf{r}^{\mathcal{R}}}{\|\mathbf{r}^{\mathcal{R}}\|_2^3} \right) + \delta\mathbf{d}^{\mathcal{R}} - \dots \\ & \dots - \mathcal{I}\dot{\boldsymbol{\omega}}^{\mathcal{R}} \times \delta\mathbf{r}^{\mathcal{R}} - 2\mathcal{I}\boldsymbol{\omega}^{\mathcal{R}} \times \delta\dot{\mathbf{r}}^{\mathcal{R}} - \mathcal{I}\boldsymbol{\omega}^{\mathcal{R}} \times \mathcal{I}\boldsymbol{\omega}^{\mathcal{R}} \times \delta\mathbf{r}^{\mathcal{R}} \end{aligned} \quad (4)$$

Here, \mathbf{r} and $\mathcal{I}\boldsymbol{\omega}^{\mathcal{R}}$ are the position and \mathcal{R} -frame angular velocity vectors of the observer with respect to an Earth-centered inertial frame (\mathcal{I}), and

$$\delta\mathbf{d}^{\mathcal{R}} \triangleq \mathbf{d}_t^{\mathcal{R}} - \mathbf{d}_o^{\mathcal{R}} = (\delta d_R, \delta d_T, \delta d_N)^T \quad (5)$$

is the \mathcal{R} -frame differential acceleration vector, which accounts for all control inputs and perturbations affecting the spacecraft motion. Note that the time derivatives in Eq. (5) are taken in frame \mathcal{R} .

Hereafter, subscripts t and o refer to the target and observer spacecraft, respectively. It is important to note that, even in the Keplerian two-body orbital motion scenario (i.e., $\delta\mathbf{d}^{\mathcal{R}} = \mathbf{0}$), no exact closed-form solutions exist for Eq. (4). Instead, by invoking the assumption that the relative separation is small with respect to the observer absolute orbit radius, the differential gravity terms on the right side of Eq. (4) can be linearized. The resulting equations of motion are known as the Lawden or Tschauner–Hempel equations [27,28]. Although several approaches exist for solving these equations, the STM developed by Yamanaka and Ankersen [29] is considered the state-of-the-art model, and is proposed for onboard implementation in the Proba-3 mission [30]. If the additional assumption of a circular

reference orbit is enacted, the linearized equations reduce further to the well-known Hill–Clohessy–Wiltshire (HCW) equations [31], for which there exists a simple STM solution. Even though this model has strong flight heritage dating back to the Apollo era, it is largely hampered by a limited range of applicability due to linearization, preclusion of perturbations in the model, and the fast time-variation of the state elements making simple deterministic guidance and control strategies difficult to devise.

2. Relative Orbital Elements

As an alternative state, the ROEs used in this paper are made up of combinations of Keplerian orbital elements (OEs) of the observer and target spacecraft. Many options are available for the definition of this OE state, including singular elements [32], quasi-nonsingular elements [33], and nonsingular ([34] p. 25) elements. For this paper, the quasi-nonsingular vector

$$\boldsymbol{\alpha} = (a, e_x, e_y, i, \Omega, u)^T \quad (6)$$

constitutes the absolute OE vector of a spacecraft, where a is the semimajor axis, e_x and e_y are the eccentricity vector components, i is the orbit inclination, Ω is the right ascension of the ascending node (RAAN), and u is the mean argument of latitude. These OEs are only singular for strictly equatorial orbits, where Ω is undefined. When talking about OEs and ROEs, a distinction must be made between *osculating* and *mean* elements; the former describe the instantaneous perturbed trajectory, whereas the latter describe the trajectory subject to an averaged perturbation effect. To avoid confusion, the notation $\tilde{(\cdot)}$ will be used to denote an osculating quantity. The evolution of the OE state comes from the Gauss variational equations (GVEs), where the time variation of the osculating OEs due to an acceleration vector described in the local frame of the spacecraft is given by

$$\frac{d\tilde{\boldsymbol{\alpha}}}{dt} = G(\tilde{\boldsymbol{\alpha}})\mathbf{d}^{\mathcal{R}} \quad (7)$$

with $G \in \mathcal{R}^{6 \times 3}$ denoting the well-documented GVE matrix ([34] pp. 34–36). This set of absolute OE variational equations can be used now to derive a set of variational equations for the osculating ROEs, $\delta\tilde{\boldsymbol{\alpha}}$. First, let the ROEs be defined functionally as some linear or nonlinear combination of absolute OEs $\delta\tilde{\boldsymbol{\alpha}} = \mathbf{g}(\tilde{\boldsymbol{\alpha}}_o, \tilde{\boldsymbol{\alpha}}_t)$. Then, the time evolution of the ROEs simply comes from differentiating the vector-valued function \mathbf{g} with respect to time and substituting the individual contributions from the GVEs for each spacecraft, yielding

$$\begin{aligned} \frac{d\delta\tilde{\boldsymbol{\alpha}}}{dt} = & \left[\frac{\partial \mathbf{g}}{\partial \tilde{\boldsymbol{\alpha}}_o} G(\tilde{\boldsymbol{\alpha}}_o) + \frac{\partial \mathbf{g}}{\partial \tilde{\boldsymbol{\alpha}}_t} G(\tilde{\boldsymbol{\alpha}}_t)^{\mathcal{R}_t} \mathcal{R}^{\mathcal{R}} \right] \mathbf{d}_o^{\mathcal{R}} + \dots \\ & \dots + \left[\frac{\partial \mathbf{g}}{\partial \tilde{\boldsymbol{\alpha}}_t} G(\tilde{\boldsymbol{\alpha}}_t)^{\mathcal{R}_t} \mathcal{R}^{\mathcal{R}} \right] \delta\mathbf{d}^{\mathcal{R}} \triangleq \mathbf{f}(\delta\tilde{\boldsymbol{\alpha}}, \tilde{\boldsymbol{\alpha}}_o, \mathbf{u}, \mathbf{p}, t) \end{aligned} \quad (8)$$

The rotation matrix $\mathcal{R}_t \mathcal{R}^{\mathcal{R}}$ transforms $\mathbf{d}_t^{\mathcal{R}}$ from \mathcal{R} to the targets RTN frame, \mathcal{R}_t , in order to apply the GVEs. Notice from Eq. (8) that the total ROE evolution is intuitively composed of variations due to the perturbation of the reference orbit and variations due to differential accelerations in the observer–target system. Although Eqs. (4) and (8) fundamentally describe the same underlying physics, the effectiveness of the ROE state is evident when considering again the Keplerian two-body example. Unlike Eq. (4), the solution of Eq. (8) in this case is trivial; for orbits of equal energy the ROEs are all constant. Furthermore, under the effect of perturbations, the osculating ROEs evolve slowly in time with respect to the reference orbital period and can be decomposed into short-period (on the order of the orbital period) oscillations, long-period (an order of magnitude larger than the orbital period) oscillations, and secular variations. This slow time variation allows larger time steps to be used when numerically integrating the GVEs, effectively increasing propagation efficiency for comparable modeling accuracy. Instead, averaging theory can be applied to obtain the variational equations for the mean ROEs subject to the secular perturbation effects. This method leads to tractable

equations of motion that can be solved in closed-form, as demonstrated by both Koenig et al. [23] for J_2 and differential atmospheric drag perturbations, and Guffanti et al. [35] and Guffanti and D'Amico [36] for solar radiation pressure (SRP) and lunisolar third-body perturbations. Within an angles-only navigation filter, these models enable accurate and efficient modeling of the environment in LEO and GEO, respectively, and their combination is used to bridge the gap between low- and high-altitude phases of HEO scenarios.

In general, the choice of ROE function \mathbf{g} is dependent on the orbital scenario. For this paper, the ROEs are defined as the following combination of quasi-nonsingular absolute OEs:

$$\delta \mathbf{e} = \begin{pmatrix} \delta a \\ \delta \lambda \\ \delta e_x \\ \delta e_y \\ \delta i_x \\ \delta i_y \end{pmatrix} = \begin{pmatrix} \delta a \\ \delta \lambda \\ \|\delta \mathbf{e}\|_2 c_\varphi \\ \|\delta \mathbf{e}\|_2 s_\varphi \\ \|\delta \mathbf{i}\|_2 c_\vartheta \\ \|\delta \mathbf{i}\|_2 s_\vartheta \end{pmatrix} \triangleq \begin{pmatrix} (a_t - a_o)/a_o \\ u_t - u_o + (\Omega_t - \Omega_o)c_{i_o} \\ e_{x,t} - e_{x,o} \\ e_{y,t} - e_{y,o} \\ i_t - i_o \\ (\Omega_t - \Omega_o)s_{i_o} \end{pmatrix} \quad (9)$$

Specifically, δa is the relative semimajor axis, $\delta \lambda$ is the relative mean longitude, $\delta \mathbf{e}$ is the relative eccentricity vector, and $\delta \mathbf{i}$ is the relative inclination vector. The notations c and s denote the cosine and sine of the angle in the subscript. In the second form of the ROEs in Eq. (9), the relative eccentricity and inclination vectors have been written in polar form in terms of their magnitudes and phase angles, φ and ϑ , respectively. The former angle represents the relative periaapsis, whereas the latter represents the relative ascending node. Just as with the absolute OEs defined in Eq. (6), this definition of the ROEs is only singular for strictly equatorial orbits, and is equally applicable to osculating and mean elements.

This definition of mean quasi-nonsingular ROEs in particular has an insightful geometric interpretation in near-circular orbits wherein each element corresponds to distinct features of the relative trajectory in the RTN frame. Specifically, δa and $\delta \lambda$ capture mean offsets in the radial and along-track directions, respectively; the magnitudes of $\delta \mathbf{e}$ and $\delta \mathbf{i}$ correspond, respectively, to the magnitudes of the in-plane (radial/along-track) oscillations and out-of-plane (cross-track) oscillations; and their phase angles φ and ϑ dictate the orientation and warping of the tilted ellipse in the radial/cross-track plane [13] (see Fig. 1a). In the context of far-range angles-only navigation, $\delta \lambda$ is a strong approximation to the weakly observable range. D'Amico [33], D'Amico and Montenbruck [37] used these features to formulate the so-called relative eccentricity/inclination-vector separation guidance laws, which are useful for establishing passively safe relative trajectories in the presence of along-track position uncertainty. Koenig and D'Amico later generalized the relative eccentricity/inclination-vector separation approach from binary formations to arbitrarily-sized swarms [38]. The near-circular orbit geometric interpretation was extended to arbitrarily eccentric orbits by Sullivan and D'Amico [22] with the introduction of a modified combination of the mean quasi-nonsingular ROEs denoted by $(\delta a, \delta \lambda^*, \delta \mathbf{e}^*, \delta \mathbf{i})^T$. In addition to the aforementioned geometric features, the new interpretation superimposes oscillations in the radial and along-track directions that are proportional to the observer orbit eccentricity e_o . Figure 1b shows the dimensionless RTN frame geometry, with components proportional to e_o shown in red.

B. Angles-Only Measurement Model

The angles-only measurement model describes the relationship between the internal dynamical state used in the filter and the measurements received by the VBS. As previously mentioned, this work uses an internal state composed of the quasi-nonsingular ROEs defined in Eq. (9), which is augmented to include additional parameters as described in Sec. III. The VBS measurements at each sample time are a pair of synchronous bearing angles, denoted as the azimuth, α , and the elevation, ϵ . These bearing angles can be interpreted in several ways. From a purely geometric standpoint, the bearing angles subtend a line-of-sight (LOS) vector in the observer's VBS frame, \mathcal{V} ,

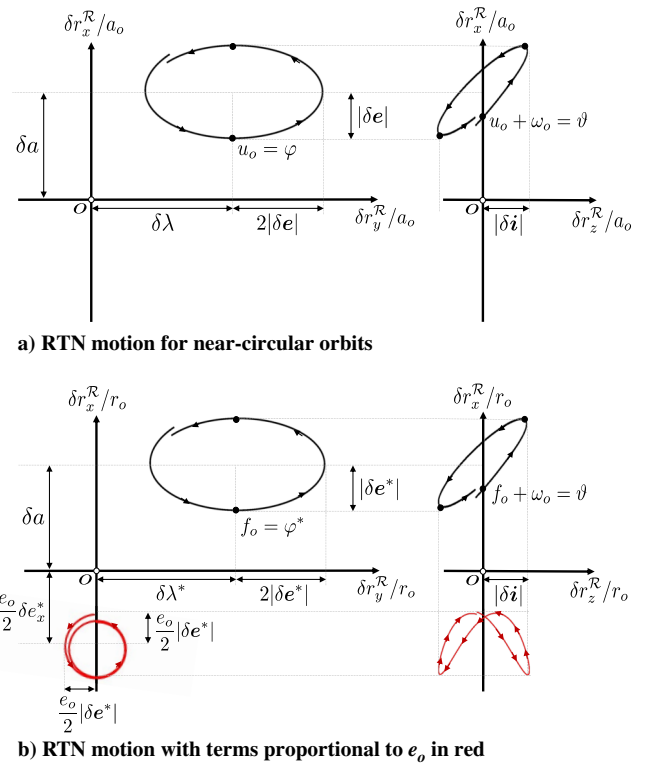


Fig. 1 Nondimensional RTN geometric interpretation of ROEs for a) near-circular and b) eccentric orbits.

which points from the observer to the target. Instead, one might also consider these angles as an analog for the image frame coordinates to the target object pixel cluster. Figure 2 highlights the former interpretation, where the \mathcal{V} and \mathcal{R} frames are shown, and the bearing angles are defined with respect to the VBS-frame relative position vector, $\delta \mathbf{r}^{\mathcal{V}}$. Formally, the bearing angle measurement model is written with respect to $\delta \mathbf{r}^{\mathcal{V}}$ as

$$\mathbf{y} = \mathbf{h}'(\delta \mathbf{r}^{\mathcal{V}}) = \begin{pmatrix} \alpha \\ \epsilon \end{pmatrix} = \begin{pmatrix} \arcsin(\delta r_y^{\mathcal{V}} / \|\delta \mathbf{r}^{\mathcal{V}}\|_2) \\ \arctan(\delta r_x^{\mathcal{V}} / \delta r_z^{\mathcal{V}}) \end{pmatrix} \quad (10)$$

Clearly, Eq. (10) is only one part of the complete measurement model \mathbf{h} . The algorithm for transforming from mean or osculating ROEs to bearing angles is highlighted in Algorithm 1. There are two particularly important components to this model that deserve attention. First, when using mean elements it is important to convert to osculating elements if possible (see Algorithm 1, lines 4 and 5). This ensures that the bearing angles, which live in an inherently "oscillating" space, are matched to properly oscillating ROE trends. Schaub




Fig. 2 Bearing angle geometry in the vision-based sensor frame, \mathcal{V} .

Algorithm 1: Complete measurement model

```

1: function  $h(\delta\mathbf{e}, \mathbf{e}_o, {}^V R^T)$ 
2:    $\mathbf{e}_r$  from Eq. (9) using  $\delta\mathbf{e}$  and  $\mathbf{e}_o$ 
3:   if using mean elements then
4:      $\mathbf{r}_o^T \leftarrow \text{MEANToOSC}(\mathbf{e}_o)$ 
5:      $\mathbf{r}_r^T \leftarrow \text{MEANToOSC}(\mathbf{e}_r)$ 
6:   else
7:      $\mathbf{r}_o^T \leftarrow \mathbf{e}_o$ 
8:      $\mathbf{r}_r^T \leftarrow \mathbf{e}_r$ 
9:   end if
10:   $\delta\mathbf{r}^V \leftarrow {}^V R^T(\mathbf{r}_r^T - \mathbf{r}_o^T)$ 
11:   $(\alpha, \epsilon) \leftarrow \delta\mathbf{r}^V$  from Eq. (10)
12:  RETURN  $(\alpha, \epsilon)$ 
13: end function

```

and Junkins [39] and Alfriend ([34] Appendix F) both provide an analytical model for the mean-to-osculating transformation that is linear with respect to J_2 , whereas more recent work by Gaias et al. [40] has put forth additional transformations to account for the remaining terms of the geopotential. Instead, there are limited closed-form models for these transformations that capture the effects of other perturbations; in these cases, the transformation is simply identity. Second, in order to compute the relative position vector (or LOS) vector, it is first necessary to transform the absolute OEs to absolute inertial position and velocity components, as in Algorithm 1, lines 4 and 5 or lines 7 and 8. This nonlinear transformation is provided by Schaub and Junkins [39], and retains the curvature of the trajectories inherently captured by the OE description. The measurement model obviously contains several distinct sources of useful nonlinearities that can be used to disambiguate the weakly observable interspacecraft range [22]; a major motivation for using the UKF in the sequential estimation approach stems from the desire to retain these nonlinearities by forgoing Taylor series linearization of the measurement and/or dynamics models.

As a last consideration in this discussion, notice from Algorithm 1 that knowledge of the absolute orbit and attitude of the observer is required. One possible controlled attitude configuration is shown in Fig. 2. In this orientation, the VBS camera boresight, $\hat{\mathbf{z}}^V$, points along the anti-flight direction, and the $\hat{\mathbf{y}}^V$ vector is aligned normal to the plane of the observer's orbit.

C. High-Fidelity Numerical Simulation

The numerical propagation of orbital states can be used to provide a rigorous "ground truth" model against which to compare analytical closed-form solutions. Although the computational simplicity of closed-form solutions might be desirable for onboard implementation in an estimator, numerical integration of the equations of motion can alternatively be used for state propagation when increased accuracy is required. Accordingly, this section highlights some of the main principles of numerical orbit trajectory modeling with the intent of developing a ground truth to evaluate the aforementioned analytical models, as well as to motivate numerical propagation for use in angles-only navigation.

The trajectories of Earth-orbiting satellites are influenced by multiple perturbing accelerations of various magnitudes. Montenbruck and Gill [41] provide an extensive mathematical description of these perturbations with a keen focus on modeling for high-fidelity orbit propagation. With the particular propagator used in this paper, the accelerations due to nonspherical gravitation are accounted for using a high-order/degree (up to 120×120) spherical expansion of the GGM01S Grace Gravity Model [1] for the Earth geopotential. Lunisolar third-body point-mass gravitation effects are simulated using analytic ephemeris models for the sun and moon. The acceleration due to atmospheric drag is treated with a standard V^2 law, using the NRLMSISE-00 [42] atmospheric density model. Finally, SRP perturbations use a cannonball model assuming conical Earth shadowing.

It is clear from the discussion above that the satellite orbital trajectories can be modeled using a variety of state descriptions. Just as in the case of closed-form solutions, the selection of state representation and corresponding equations of motion plays an important role in numerical propagation efficacy. While the equations of motion in translational elements (the fundamental orbital differential equations, FODEs) or OEs (the GVEs) describe the same underlying physics, the quality of performance in their numerical integration can vary substantially because of the intrinsic difference between the state representations. For example, the inertial translational elements are rapidly varying even with no applied perturbations, whereas the OEs are constant for no perturbations and slowly time-varying when perturbed. It stands to reason then that, in order to obtain equivalent propagation precision, numerically integrating the translational element equations of motion would require a smaller time step than that used with the GVEs.

As evidence for the GVE propagation efficiency, a numerical integration assessment has been conducted to study the performance of the two classes of dynamics models as a function of integration step size. Precise orbit determination products from the PRISMA mission [3] act as the ground truth data by which the outputs of the numerical propagators are evaluated. For the OE case, the absolute GVEs in Eq. (7) are numerically integrated; for the translational element case, the absolute FODEs are numerically integrated. In both cases, high-fidelity models of Earth geopotential, atmospheric drag, SRP, and lunisolar third-body perturbations are used (see Table 1). Using 200 evenly spaced initial condition sets from within the flight data, the orbits are propagated with integration step size varied in the set $\Delta t \in \{10, 30, 60, 120\}$ s. The mean 3D root-mean-squared (RMS) position propagation errors and $1-\sigma$ standard deviations are computed from the 200 sample trajectories at each time step and shown in Fig. 3; the standard deviations are captured by the shaded regions about the solid mean error trends. It is important to note that similar trends to those depicted are also found for the velocity propagation errors. Clearly, numerically integrating the GVEs allows for superior propagation accuracy while using a step size that is an order of magnitude greater than the FODE-based test. In the angles-only navigation scenario, where measurements come at the sparse rate of 1–10 min, a larger integration step size means fewer overall calls to the state update algorithm and a corresponding increase in filter propagation efficiency.

D. Simulation-Based Algorithm Verification

To verify the navigation algorithms put forth in the following sections, a simulation pipeline is established that makes use of the aforementioned full-force orbit propagator to compute the true trajectories of the observer and target satellites. The truth-side data come from numerically integrating the GVEs subject to the perturbations outlined in the top half of Table 1. The resulting trajectories are then used to generate measurements for the algorithms in one of two ways. First, they are used as inputs into a synthetic measurement emulation model that delivers bearing angles, observer absolute orbit, and observer absolute attitude inputs that have been corrupted with noise that is consistent with commercially available sensors. Table 2 captures the various noise sources that have been injected into the simulation. In particular, the observer orbit knowledge errors are conservative values for a position, velocity, and time (PVT) solution using the Global

Table 1 Orbit propagator and PRISMA flight product details

Propagator configuration	
Integrator	Runge–Kutta Dormand–Prince
Geopotential	GRACE Gravity Model 01S 120×120
Atmospheric density	NRLMSISE-00 Model
Solar radiation pressure	Flat plate with conical Earth shadow
Third-body gravity	Lunar and solar point masses
Lunisolar ephemerides	Analytic
PRISMA flight data	
Initial epoch	March 17, 2011 00 ^h :00 ^m :00 ^s
Duration	06 (h)

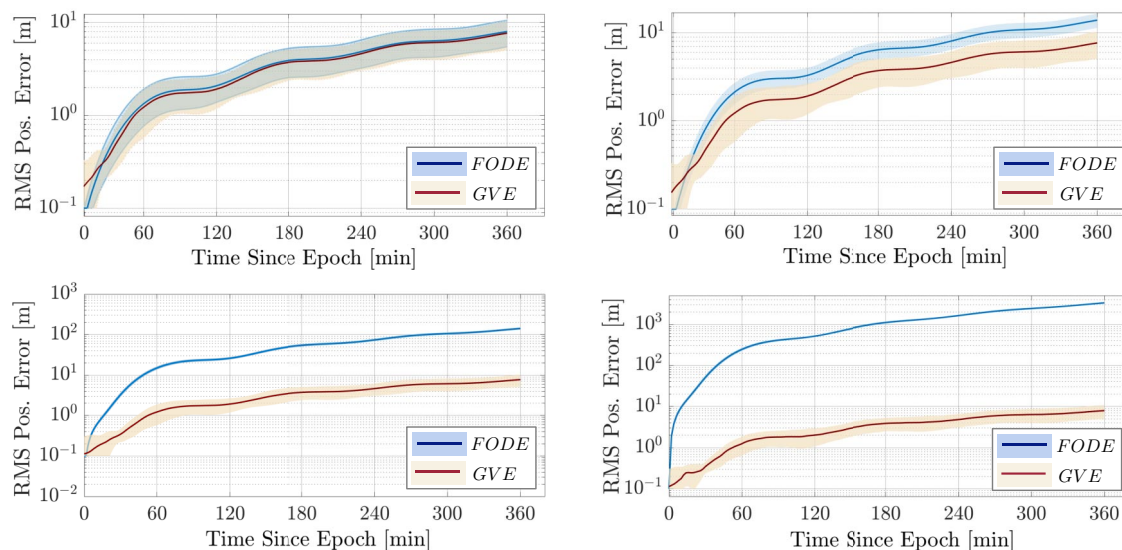


Fig. 3 Propagation performance based on choice of equations of motion for integration step sizes of 10 s (upper left), 30 s (upper right), 60 s (bottom left), and 120 s (bottom right).

Table 2 Sources of noise injected during algorithm verification

Noise source	1- σ	Unit
Observer absolute position, ^a σ_p	25	m
Observer absolute velocity, ^a σ_v	0.25	m/s
Observer boresight attitude	30	arcsec
Observer off-axis attitude	6	arcsec
(α, ϵ) measurements, σ_α and σ_ϵ	30	arcsec

^aDoubled for HEO and GEO test cases.

Navigation Satellite System (GNSS); the attitude knowledge errors are in line with commercially available attitude determination system capability; and the bearing angle noise value is consistent with results from several sensor benchtests. With regard to the latter parameter, 30 arcsec is on the order of one-third to one-half a pixel for several candidate VBS. Note that timing errors/delays are not considered yet in this work. Alternatively, the trajectories are used as inputs into a hardware-in-the-loop optical testbed [43,44] in order to stimulate a space-grade camera with dynamically rendered scenes of the observer and the background starfield. The testbed, which will be discussed in Sec. IV, delivers measurements of the bearing angles and the observer's absolute attitude to the angles-only navigation algorithms. Accordingly, only the observer absolute orbit knowledge is generated via simulation and corrupted with noise per Table 2.

Verification of the algorithms is conducted using scenarios that are composed of a primary observer absolute orbit and target relative motion configurations about that observer orbit. In this work, the primary observer orbit is parameterized by a set of mean OEs, and

categorized as one of the following from Table 3: a low-altitude (350 km) near-circular LEO (denoted *LEO1*), a near-circular sun-synchronous LEO at 850 km altitude (denoted *LEO2*), an inclined HEO (denoted *HEO1*), or a very near-equatorial geosynchronous orbit (denoted *GEO1*). The motion of targets about the primary observer orbit is instead parameterized using a set of mean ROEs. The *GEO1* scenario is near-equatorial to allow the use of the quasi-nonsingular OEs and ROEs defined previously. If operation in a strictly equatorial orbit is necessary, the nonsingular OE and ROE states defined by Alfriend ([34] p. 25) and Koenig et al. [23] can instead be applied.

For most of the test cases presented in this paper, Monte Carlo simulations are conducted that vary the target ROEs uniformly to generalize the testing over a range of reasonable formation configurations. The lower and upper bounds, \mathcal{U}_{lb} and \mathcal{U}_{ub} , for the uniform distribution that each ROE component is drawn from are shown in the lower portion of Table 3. Unless specified otherwise, the number of Monte Carlo samples is chosen as 600 to provide 95% confidence that steady-state estimation errors are consistent to within 2% of the true value for all simulation cases. Lastly, a nominal ROE configuration (denoted *RO1*) is specified for tests that do not vary the relative motion configuration. *RO1* is considered a mid-range approach configuration starting at 30 km separation with 500 m projected circular motion in the radial/cross-track plane and a mean semimajor axis offset that induces a drifting approach that brings the observer and target closer at a rate of approximately 1.5 km per orbit.

III. Angles-Only Relative Navigation Algorithms

The objective of this section is to expand upon the algorithmic features of the proposed angles-only navigation filter, and to provide

Table 3 Initial conditions for sequential algorithm verification scenarios (perigee altitude is denoted by h_p)

Observer	h_p (km)	e_x	e_y	i (deg)	Ω (deg)	u_0 (deg)
<i>LEO1</i> : low altitude	350	-5.00×10^{-4}	8.66×10^{-4}	98.70	28.87	300.00
<i>LEO2</i> : sun-synch., LTAN 11	850	-5.00×10^{-4}	8.66×10^{-4}	98.70	28.87	300.00
<i>HEO1</i> : highly elliptical	850	-2.50×10^{-1}	4.33×10^{-1}	30.00	60.00	300.00
<i>GEO1</i> : geosynchronous	35,786	-5.00×10^{-4}	8.66×10^{-4}	1.00	60.00	300.00
Relative orbit (RO)	$a\delta a$ (m)	$a\delta l$ (km)	$a\ \delta e\ _2$ (km)	φ (deg)	$a\ \delta i\ _2$ (km)	ϑ (deg)
<i>RO1</i> : midrange approach	-150	-30	0.5	0	0.5	180
Monte Carlo: \mathcal{U}_{lb}	-150	-50	0	0	0	0
Monte Carlo: \mathcal{U}_{ub}	150	-5	1	360	1	360

verification results that highlight the applicability and strengths of the new approach. Although the UKF [45] (the foundation for this estimation architecture) is no longer itself a novelty, the manner by which it is applied to strategically suit the needs of angles-only navigation is indeed a major contribution of this work. Key filter design paradigms involving the choice between linear or nonlinear dynamics modeling, the ability to adaptively tune process noise statistics in real-time, and the flexibility to estimate additional parameters of utility will all be discussed in the following.

A. Unscented Kalman Filter Dynamics Modeling

1. Linear and Nonlinear State Propagation

When a new measurement set is received, the first step in the general Kalman filtering architecture is to propagate the current state and covariance estimates forward to the measurement time. The explicit assumption within this framework is that the state dynamics are Gaussian, with the distribution mean capturing the current state estimate and the distribution covariance capturing the current state uncertainty. Although the state propagation for a nonlinear system may proceed directly through the nonlinear dynamics, the uncertainty must be updated linearly to retain a Gaussian distribution. Rather than computing a first-order Taylor series expansions of the dynamics model, the UKF enables higher-order nonlinearities to be captured by using a stochastic weighted linear regression [46]. With this insight, a navigation filter is designed that updates the state and covariance via numerical integration of the nonlinear osculating ROE dynamics given by the GVEs in Eq. (8). In conjunction with the nonlinear measurement model, this filter class (denoted hereafter as the *numerical* filter) mitigates the inherent observability problems in the angles-only navigation scenario by retaining key higher-order dynamical system features that potentially uniquely define the relative state trajectory. As an alternative to using numerical integration within the estimator, this work poses a second filtering class (denoted as the *analytic* filter) that uses an STM to linearly propagate a mean ROE estimated state and covariance. As explored in previous work [22], this method still retains some observability-improving nonlinear features that are exploited through the measurement relationship between mean ROEs and the bearing angles. The choice between a numerical or an analytic filter implementation is a strategic design tradeoff, which is quantified in the subsequent sections. Figure 4 graphically compares the two filter variants with a sampling-based example, where \mathcal{X} denotes the state

“sigma points” that are propagated through the nonlinear dynamics in the UKF and used to compute an updated empirical state covariance.

2. Effect of Orbital Perturbations

The effectiveness of either the analytic or numerical filter is driven by its ability to properly fuse the dynamical evolution of the estimation state with measured trends of the bearing angles. Because this work is intended to generalize the applicability of angles-only navigation to all relevant Earth orbit regimes, it is essential to capture the dominating perturbation effects that drive substantial variation in the ROEs over time. For the analytic filter variant, this capability hinges on the availability of STM solutions that account for perturbations in closed form. As previously mentioned, such models are formulated for the ROEs of Eq. (9) by Koenig et al. [23], Guffanti et al. [35], and Guffanti and D’Amico [36]. Instead, for the numerical filter variant, the same high-fidelity force modeling that was discussed in Sec. II can be used in the state propagation machinery for the estimator. It is important to note, however, that modeling of nonconservative perturbations like atmospheric drag and SRP requires that parameters related to the ballistic properties of the constituent satellites be known or estimated in both the linear and nonlinear models. For the nonlinear models, the ballistic properties are often accounted for through the drag and SRP ballistic coefficients. In the relative motion context, the *differential* ballistic coefficient is the primary parameter of interest for nonconservative perturbation effects and it is defined as

$$\Delta B \triangleq \frac{B_t - B_o}{B_o} \quad \text{where } B \triangleq C \frac{A}{m} \quad (11)$$

Here, C can refer either to the drag or reflectivity coefficient of the satellite, depending on whether drag or SRP is considered, and A/m is the area-to-mass ratio. The aforementioned linear dynamics models provide the ability to capture nonconservative perturbations either in terms of these same differential ballistic coefficients, or in a “model-free” sense by augmenting the dynamical state with the time rate of change of a subset of ROEs to capture the empirical average effect of the perturbation.

To depict the effect of perturbations on the ROE trajectory, a simple numerical simulation of the GVEs is conducted using the high-fidelity propagator described previously. The outputs of this qualitative study illustrate the secular, long-periodic, and short-periodic osculating ROE trends induced by the dominant Earth oblateness (J_2) perturbation in LEO and SRP in GEO. Although this represents only a subset of the perturbations that can be modeled in both the analytic and numerical filter versions, these effects are generally most dominant and effectively exemplify the types of variations expected. Furthermore, the results of these numerical simulations confirm the trends analytically predicted by Koenig et al. [23], Guffanti et al. [35], and Guffanti and D’Amico [36]. As shown in Fig. 5a, J_2 causes a drift in $\delta\lambda$ and a circular precession of the relative eccentricity vector δe . In addition to these secular and long-periodic variations, the osculating ROEs display short-period oscillations with amplitudes that grow with increasing interspacecraft separation (i.e., increasing $|\delta\lambda|$). Note that the effects of J_2 on δi are not shown in this figure. In GEO scenarios, the effect of SRP on the observer–target system is largely driven by their difference in ballistic coefficients. Figure 5b highlights the ROE evolutions due to SRP for a ΔB of 20% and 80%. In particular, drift and short-periodic oscillations in $\delta\lambda$ are accompanied by short-periodic oscillations in δa and a long-periodic circular precession of δe . The magnitude of each of these effects scales with ΔB . SRP does not generally produce substantial variation in δi .

The treatment of filter dynamics modeling as a strategic design parameter is now illustrated with a comparative case study. Both the analytic and numerical filter are implemented for the *GEO1 ROI* configuration highlighted in Table 3, and the impact of SRP is considered by accounting for or omitting it from the “onboard” modeling. In the analytic filters, the STM developed by Guffanti et al. [35] is used to capture the mean perturbation trends; in the numerical filters, a flat plate model with analytic sun ephemerides and canonical Earth shadowing is implemented to account for the osculating

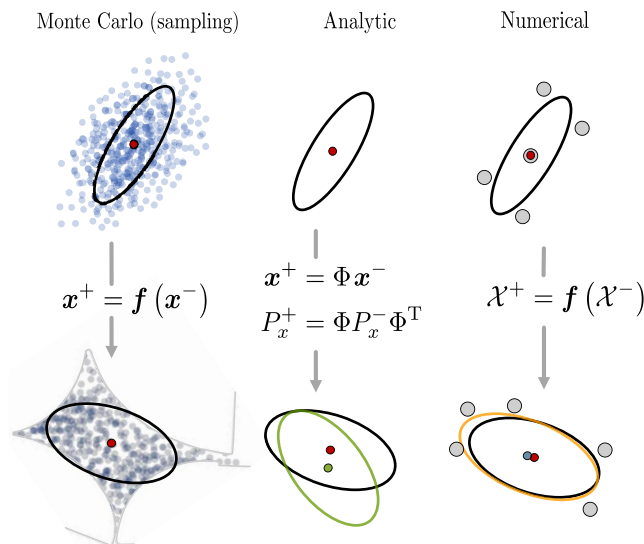




Fig. 4 Comparison of Gaussian propagation methods. The leftmost subfigure illustrates a pure sampling-based strategy similar to Monte Carlo analysis. The central illustration shows the propagation of the mean and covariance through a linear transformation. The rightmost subfigure depicts the sigma point propagation and ensuring notional mean and covariance computation from their updated distribution.



a) Effects of J_2 in LEO over 30 days




b) Effects of SRP in GEO over 60 days. $\Delta B = 20\%$ (dark) and $\Delta B = 80\%$ (light) are both shown

Fig. 5 Mean (dashed) and osculating (solid) perturbed ROE trends with initial (\blacktriangle) and final conditions (\blacksquare).


perturbation effects. In each case, the filter is given a modest initial-ization error of approximately 15% for each ROE, and receives measurements at an interval of 120 s. When accounting for SRP, it is assumed that the true differential ballistic coefficient of 80% is known exactly. Resulting error plots are shown in Fig. 6a for the analytic filter, and Fig. 6b for the numerical filter. Note that only the relative mean longitude error is shown, because it is well established that the error in all other nonzero ROE components scales with this component as a result of the observability decoupling [12,22]. The test results here strongly indicate the necessity of modeling the SRP perturbation for filter accuracy. Both filter variants display a large but convergent steady-state error and covariance in the case where SRP is omitted. The convergence to a largely biased estimate is a strong indicator of dynamics mismodeling in the filter leading to new measurements being associated to improper estimates of range. Instead, including SRP in the filter dynamics model is shown to stabilize the filter to an error well within 1%. Interestingly, because the mean dynamics model in the analytic filter does not immediately account for the short-period oscillations seen in $a\Delta\lambda$ (as seen in Fig. 5b), the estimation error in Fig. 6a is oscillatory about a mean of approximately 2.3 km. The numerical filter does not display this oscillatory response because it is integrating the osculating equations of motion.

3. Monte Carlo Testing of Sensitivity to Observer Orbit Knowledge

The filter performance benefits from using numerical integration for the state propagation do not come for free. In fact, there are two main concerns when using the numerical filter. First, nonlinear propagation of the equations of motion for even a low-fidelity force model is generally at least an order of magnitude more computationally intensive than propagating a linear dynamics model [25]. This



a) Analytic filter error and 3- σ covariance bounds



b) Numerical filter error and 3- σ covariance bounds

Fig. 6 Filter $a\Delta\lambda$ estimation error mean/covariance, excluding SRP (red/tan) and including SRP (blue) for the GEO1 ROI case.

can be particularly problematic within the UKF framework, where each sigma point must be propagated through the dynamics independently. However, as seen in the results of Fig. 3, integrating the GVEs allows for large integration time steps during the propagation process. Accordingly, the concern of inefficient filter dynamics modeling is largely allayed by working in the ROE framework.

The second concern with using the numerical filter variant stems from the reliance upon observer orbit information, which may have large errors. Propagation of the relative motion dynamics always requires some knowledge of the absolute reference orbit (in this case, the observer). This can be as minimal as knowing only the reference orbit mean motion for the case of strictly Keplerian orbits, or as complex as requiring the complete orbital state as well as ballistic properties. Without providing a metric for the uncertainty in these essential parameters (e.g., as in a consider-covariance formulation), the filter has limited means to correct the estimation trends to account for large errors in the absolute orbit state. This can lead to reduced filter performance in scenarios where the probabilistic description of the observer orbit is not well-characterized. The performance impacts due to this lack of process information can also be more or less pronounced based on how the information is coupled to the main filter state dynamics. In a nonlinear relative motion dynamics model, the knowledge of the observer's orbital parameters is tightly coupled with the propagation accuracy of the relative state. Instead, in a linear relative motion propagation method, by the very nature of linearizing the equations of motion, the practical coupling of observer state knowledge is reduced to a first-order effect.

To better understand the practical implications of this coupling of relative motion propagation and observer absolute orbit knowledge, two Monte Carlo series of 600 tests are conducted using both the analytic and numerical filters in the LEO2 observer orbit case. The numerical filter uses a 120 s time step for the internal integration of the GVEs, with force modeling to include a 5×5 Earth gravity model and lunisolar third-body gravitation. The analytic filter instead propagates a state of mean ROEs subject to only J_2 perturbations using the model from Koenig et al. [23]. As in all simulation results presented in this work, the error sources captured in Table 2 are used to corrupt "true" orbit, attitude, and sensor measurements. However, for this

series of tests, the statistics of the observer orbit error distribution are themselves sampled from a statistical distribution for each Monte Carlo iteration. For example, the j th test in this Monte Carlo series will include observer orbit position and velocity component errors ($E_{p,v}^{(j)}$) that are sample from the zero-mean Gaussian distribution

$$E_{p,v}^{(j)} \sim \mathcal{N}(0, \sigma_{p,v}^{(j)}) \quad (12)$$

However, rather than $\sigma_{p,v}^{(j)}$ being a constant (as specified in Table 2), it is sampled from the Gaussian distribution

$$\sigma_{p,v}^{(j)} \sim \mathcal{N}(\mu_{\sigma_{p,v}}, \Sigma_{\sigma_{p,v}}) \quad (13)$$

Two separate Monte Carlo series are conducted, and are delineated based on their values of $\mu_{\sigma_{p,v}}$ and $\Sigma_{\sigma_{p,v}}$ as shown in Table 4. Additionally, for each Monte Carlo iteration, the true ROE state defining the relative motion is sampled from the uniform distribution provided in Table 3. The filter is initialized at each iteration with approximately 25% error with respect to the true values. Results for these tests are shown in Fig. 7, with the primary figure of merit being the percentage error in the relative mean longitude. Note that the mean of the error distributions is shown in solid lines, whereas the 1- σ empirical standard deviation bounds are shown shaded around the mean trend. For additional realism, the filter does not receive angular measurements during periods of eclipse (gray-shaded regions of the error plots).

The Monte Carlo test results show a few interesting features related to the filter sensitivity to observer absolute orbit knowledge. First in the “moderate” observer error case (Fig. 7a), both the numerical and analytic filters perform quite well with regard to percentage error in the range analog, $a\delta\lambda$; the former reaches a final error distribution of $0.9 \pm 1.3\%$, whereas the latter converges to $1.7 \pm 2.4\%$ after

five orbits without orbital maneuvering. Note, however, that the filter estimation is converged and stable after about three orbits. As expected, the numerical filter generally outperforms the analytic filter in both mean error and precision of convergence due to improved nonlinear dynamics modeling. Instead for the “conservative” orbit knowledge error case (Fig. 7b), the numerical filter fails to stabilize to an accurate estimate of $a\delta\lambda$ on average, with an error distribution of $21.1 \pm 6.4\%$ after 18 orbits and no clear convergent trend. Note that the numerical filter has a consistent tendency to overestimate the relative mean longitude in the high-noise case. Unlike the numerical filter, the analytic variant displays a clear convergent trend, reaching an error distribution of $4.7 \pm 4.1\%$ at the end of simulation. The transient error phase is more pronounced and the precision of convergence is generally larger as compared with the same filter in the “moderate” case. Still the filter generally performs successfully if provided sufficient (in this case, nearly a factor of four more) time to converge. These results reveal a sort of duality in the observability-constrained nonlinear filtering framework: high levels of noise in a tightly coupled nonlinear model can actually prove more detrimental to filter stability than neglecting observability-improving system nonlinearities in favor of noise robustness with linearized models. Within this trade space, these comparative results ultimately indicate the interplay between time to convergence, relative state estimation accuracy, stability, and robustness, which must be evaluated against the accuracy in observer orbit that is expected for the particular mission scenario.

B. Adaptive Process Noise Tuning

Now that it has been firmly established that the observability constraints imposed by angles-only navigation require reasonable state propagation accuracy for effective estimation performance, this section highlights a method of supplementing the onboard dynamics model using the filter process noise covariance. In particular, this approach uses sequences of bearing angle measurement residuals to inform how the underlying state uncertainty should be evolving. As a result, the propagation of the formal state covariance is enhanced by using measurement information, which evolves under the true dynamics, in cases where the filter employs a reduced-dynamics propagation method.

This work primarily uses the method of covariance-matching for adaptive process noise tuning that was introduced by Myers and Tapley [47,48] for general orbit determination, and later expanded upon for satellite formation-flying by Fraser [49]. A key trait of this approach is the use of a sliding window of measurement residuals whose empirical covariance is related to the process noise covariance through the Kalman filter gain matrix. In terms of the STM and state formal covariance matrix, Σ , the process noise covariance matrix at time t_k can be adaptively updated according to

$$Q_k = \frac{1}{N} \sum_{i=k-N}^{k-1} (\Sigma_{i|i} - \Phi_{i,i-1} \Sigma_{i-1|i-1} \Phi_{i,i-1}^T + \Delta_i^x \Delta_i^{xT}) \quad (14)$$


Here, N is the length of the sliding window, and the state residuals, Δ_i^x , in terms of the measurement residual and Kalman gain, \mathcal{K} , are

$$\Delta_i^x = \mathcal{K}_i (y_i - \mathbf{h}(\mathbf{x}_{i|i-1})) = \mathcal{K}_i \Delta_i^y \quad (15)$$


Note that a subscript $i|i-1$ indicates an estimate at time t_i using measurements up to time t_{i-1} . Following the covariance update, two additional steps are taken to complete the process noise tuning. First, a key advantage of the tuning law in Eq. (14) is that it does not require invoking an assumption of steady-state filter performance and can be equally applied during transient phases of estimation; however, this comes with the added complexity that Q_k is not guaranteed to be positive semidefinite (PSD), and is therefore not immediately valid as a covariance matrix. Accordingly, this work draws from the body of work on least-squares covariance adjustment problems (LSCAPs) [50] to formulate an optimal transformation of the process noise matrix estimate as

Table 4 Monte Carlo sampling statistics for the observer orbit error

Error case	μ_{σ_p} (m)	$\sqrt{\Sigma_{\sigma_p}}$ (m)	μ_{σ_v} (m/s)	$\sqrt{\Sigma_{\sigma_v}}$ (m/s)
Moderate	15	3	0.15	0.03
Conservative	500	100	5	1



a) Estimation error for “moderate” observer error case



b) Estimation error for “conservative” observer error case

Fig. 7 Monte Carlo filter performance for absolute orbit knowledge sensitivity.

$$\begin{aligned} & \text{minimize} && (1/2)\|\hat{Q}_k - Q_k\|_F^2 \\ & \text{subject to} && \hat{Q}_k \geq 0 \end{aligned} \quad (16)$$

Higham [51] proposes a solution to this LSCAP, which takes the form

$$\hat{Q}_k = (1/2)(B + \mathcal{H}) \quad (17)$$

where \mathcal{H} is a symmetric polar factor of $B = (1/2)(Q_k + Q_k^T)$. Algorithm 2 highlights the singular value decomposition approach used in this paper to numerically compute the optimal covariance adjustment. Second, because the inclusion of adaptive tuning intrinsically provides another channel by which the estimator can diverge, a regularization feature is included to ensure that the process noise covariance matrix does not grow excessively large over short time spans. Algorithm 2 employs a heuristic regularization based on the spectral radius, $\rho(Q)$, of the successive process noise covariance matrices:

$$\frac{\rho(\hat{Q}_k)}{\rho(\hat{Q}_{k-1})} \leq \bar{\rho} \quad (18)$$

for some user-specified maximum allowable spectral radius ratio, $\bar{\rho}$.

A fitting illustration of the adaptive filter effectiveness in supplementing onboard dynamics modeling for improved estimation performance comes in the *HEOI* test case. Here, the observer and target are in HEOs, traversing regions with high SRP influence (at high altitude) and regions with high J_2 influence (at lower altitudes). For this case study, both the standard analytic and adaptive analytic filters are employed using a linear dynamics model that only accounts for J_2 . The performance of both filters is assessed in Monte Carlo simulation, where the ROEs defining the formation geometry are sampled according to Table 3. Additionally, for each Monte Carlo simulation, the differential SRP ballistic coefficient for the formation is sampled out of a uniform distribution $\mathcal{U}(-100\%, 100\%)$. Figure 8 depicts the estimation errors and $1\text{-}\sigma$ empirical standard deviation bounds from the

Algorithm 2: Adaptive process noise tuning

```

1: function UPDATEQ ( $x_{1:N|1:N}, \Sigma_{1:N|1:N}, \Delta_{1:N}^y, \mathcal{K}_k, \bar{\rho}$ )
2:   Tune  $Q_k$  from Eq. (14)
3:   Least-Squares Covariance Adjustment:
4:    $B \leftarrow (1/2)(Q_k + Q_k^T)$ 
5:    $U, S, V^* \leftarrow \text{SVD}(Q_k)$ 
6:    $\mathcal{H} \leftarrow VSV^T$ 
7:    $\hat{Q}_k \leftarrow (1/2)(B + \mathcal{H})$ 
8:   Regularize  $\hat{Q}_k$  growth:
9:    $\rho_k \leftarrow \text{EIG}(\hat{Q}_k)$ 
10:   $\rho_{k-1} \leftarrow \text{EIG}(\hat{Q}_{k-1})$ 
11:  if  $\rho_k/\rho_{k-1} \geq \bar{\rho}$  then
12:     $\hat{Q}_k \leftarrow \bar{\rho} \frac{\rho_{k-1}}{\rho_k} \hat{Q}_k$ 
13:  end if
14:  RETURN  $\hat{Q}_k$ 
15: end function

```

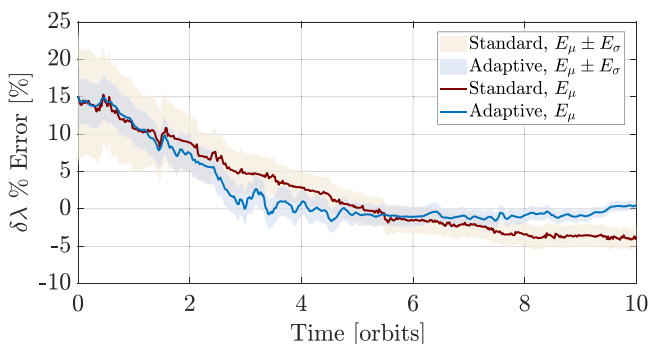


Fig. 8 Monte Carlo comparison of $\delta\lambda$ estimation error for the standard and adaptive analytic UKF in *HEOI* scenario.

simulation campaign. Compared with the standard filter, the adaptive filter demonstrates better average convergence time (about 4 orbits vs 8 orbits), lower mean estimation error at steady state (within 1% vs 5%), and tighter estimation error standard deviation (0.35% vs 1.10%). The inclusion of adaptive process noise tuning increases the computational runtime by approximately 8% compared with the non-adaptive time-update step.

Other than acting to complement reduced-dynamics state propagation in the filter, adaptive process noise tuning has another compelling outcome. Because the process noise statistics are modified as a function of measurement residual histories, the adaptive filter is more receptive to new measurements and is expected to better handle anomalous events. As an example of this filter receptiveness, consider the following case: an observer is estimating the orbital state of a non-cooperative target space object, and part way through the filter convergence phase, the target executes an unknown orbital maneuver that changes its long-term trajectory. Furthermore, due to visibility limitations and eclipses, the filter is only able to receive measurements for approximately 30% of the orbit. The filter performance for this demanding scenario in the *LEO2 ROI* case is shown in Fig. 9. Estimation error and uncertainty bounds in the relative mean longitude ROE component are shown for both the standard and adaptive forms of the analytic filter. As always, periods of measurement outage are shown in the shaded gray regions. Notice that, following the unknown target maneuver, the adaptive filter covariance (shaded blue) displays a pronounced increase, whereas the standard filter covariance increases only slightly due to the mismatch. The large change in the adaptive filter arises because the estimator has computed disproportionately erroneous measurement residuals, and the adaptive tuning component uses those to increase process noise covariance. In so doing, the filter trusts its onboard dynamics model and state estimate less, and accordingly weights incoming measurements more heavily. The result is a faster stabilization of the adaptive filter estimate as compared with the standard filter. Whereas it takes the adaptive filter only about four orbits to return convergence to within $4.8 \pm 1.6\%$ of the true $a\delta\lambda$, the standard filter remains on a slow convergent trend for the remaining seven simulated orbits. The final steady-state error of the adaptive filter is reduced down to within $0.8 \pm 0.7\%$ by the end of the simulated estimation window.

C. Estimating First-Order Gauss–Markov Processes

In addition to solving for the relative orbital state of target space objects, the ability to estimate additional relevant parameters extends the usefulness and robustness of the filter. As mentioned in prior sections, the differential ballistic coefficient for SRP or drag is a parameter that is required in either linear or nonlinear propagation models and is generally unknown at the start of non-cooperative operations. Furthermore, the differential ballistic coefficient provides insight into the relative size and attitude of the space object. Accordingly, a filter that can estimate this quantity provides a wealth of new information that improves overall situational awareness and modeling capability. As another example of additional relevant parameters, consider that angles-only navigation relies on a calibrated optical sensor to detect variations in the relative motion (at the subdegree level). The capability to estimate and calibrate sensing biases enables more robust performance and greater ability

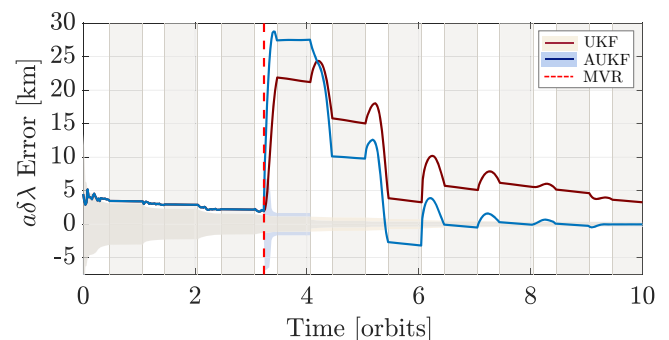


Fig. 9 Estimation error and $3\text{-}\sigma$ covariance bounds with an unknown target maneuver in *LEO2 ROI*.

to detect and isolate sensing faults. This section highlights the method by which these specific parameters can be captured in the filter through state augmentation. It is important to note that this approach can be generalized to myriad additional parameters, as was shown in earlier work by Sullivan and D’Amico [13] for estimating empirical accelerations.

The sensor biases and differential ballistic coefficient parameters can be reasonably modeled as band-limited noise processes. Accordingly, a first-order Gauss–Markov model [52] is used to capture the evolution of such a parameter, here defined generically as β :

$$\dot{\beta}(t) = -\frac{1}{\tau}\beta(t) \Rightarrow \beta(t + \Delta t) = \beta(t) \exp\left(-\frac{\Delta t}{\tau}\right) \quad (19)$$

The parameter τ is the autocorrelation time constant, which can be tuned according to the expected variability in the parameter. For a stationary bias (i.e., for an Ornstein–Uhlenbeck process), this time constant is set to ∞ to capture a random walk process. For a nonstationary parameter, the time constant is chosen to effectively sample the variation over the estimation period.

Monte Carlo simulation results for the estimation of azimuth and elevation angle biases are shown for the *HEOI ROI* case in Fig. 10. Sensor biases are treated as stationary parameters ($\tau = \infty$) and true values are sampled from a normal distribution $b_\alpha, b_e \sim \mathcal{N}(50, 25)$ arcsec for each simulation. Note that these values for mean and standard deviation correspond to approximately 2 and 1 pixels of a representative VBS, respectively. The filter estimation state is augmented to include both additive bias terms, yielding the new dynamical system

$$\begin{aligned} \mathbf{x}(t) &= [\delta(t), b_\alpha, b_e]^T \in \mathbf{R}^3 \\ \dot{\mathbf{x}}(t) &= [f(\delta(t), o(t), \mathbf{p}(t), t), 0, 0]^T \\ \mathbf{y}(t) &= \mathbf{h}(\delta(t), o(t), \nu R^T) + [b_\alpha, b_e]^T \in \mathbf{R}^2 \end{aligned} \quad (20)$$

The filter employed here is the numerical UKF variant (with unscented transform parameters $\alpha = 10^{-3}$, $\beta = 2$, $\kappa = 1$) that models a 5×5 Earth gravity model, lunisolar gravitation, and SRP (with a differential ballistic coefficient error of 5%), and is initialized with ROE error of 25% (relative to the true ROEs) and angle biases of 0 arcsec. The results indicate that the azimuth angle bias is generally estimated to high accuracy very quickly, with a strongly stable convergence to within 2 arcsec ($3\text{-}\sigma$). Instead, the elevation bias angle shows a substantially slower convergence trend, with a larger uncertainty in the steady-state error distribution. The estimation behavior for these biases makes sense when recalling the measurement model: the azimuth angle captures out-of-plane angular deviations with respect to the observer’s orbit, whereas the elevation angle captures in-plane angular deviations from the observer’s (anti-)along-track direction (see Fig. 2). These angles are generally periodic over the orbital period, and, for the case of the azimuth angle, centered on a mean angle of zero. Instead, by definition, the elevation angle can have two sources of natural

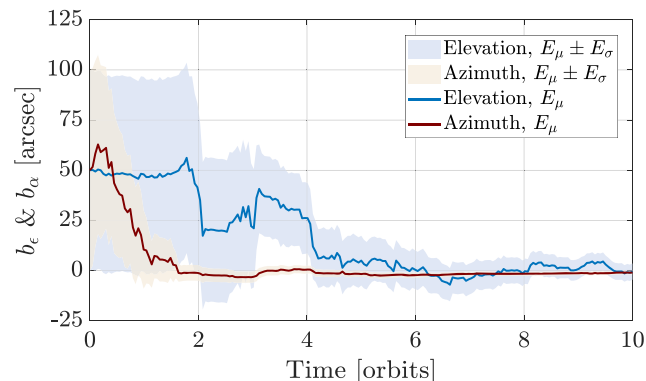


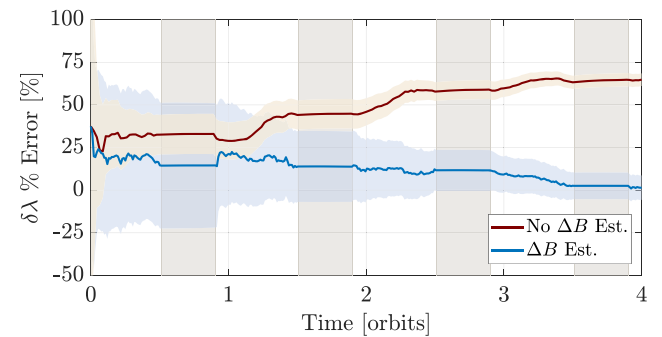
Fig. 10 Monte Carlo mean (solid) and $1\text{-}\sigma$ (shaded) estimation error in bearing angle bias for the *HEOI* test case.

nonzero mean angle: 1) a nonzero difference in semimajor axis between observer and target, and 2) a large nonzero mean along-track separation between observer and target. In the first case, there is a rectilinear radial offset caused by the target being on a higher or lower orbit; in the second case, the large along-track separation results in a curvilinear radial offset due to the orbit curvature, which, when perceived by the observer, manifests as an offset in elevation angle. Because the elevation angle is inherently composed of these two possible nonzero offsets, the estimation of any elevation angle bias (i.e., due to sensing miscalibration) is more tightly coupled to the estimation of the inter-object range and associated observability constraints. Still, the filter is able to estimate elevation angle bias to within approximately 5 ± 2 arcsec after about five orbits.


Now consider the case where the observer has very poor knowledge of the target ballistic properties for either modeling the SRP force of the atmospheric drag force. The estimation state can be augmented to solve for such a set of parameters to improve the force modeling accuracy in the filter dynamics model, yielding the dynamical system

$$\begin{aligned} \mathbf{x}(t) &= [\delta(t), \Delta B(t)]^T \in \mathbf{R}^2 \\ \dot{\mathbf{x}}(t) &= \left[f(\delta(t), o(t), \Delta B(t), \mathbf{p}(t), t), -\frac{\Delta B(t)}{\tau} \right]^T \\ \mathbf{y}(t) &= \mathbf{h}(\delta(t), o(t), \nu R^T) \in \mathbf{R}^2 \end{aligned} \quad (21)$$

Note that the ballistic coefficients play no direct role in the measurement modeling of the bearing angles, but are indeed coupled into the internal dynamics model describing the relative motion. Monte Carlo simulation results are shown in Figs. 11 and 12 for the *LEOI* and *GEOI* observer cases, respectively. The relative orbit initial conditions for the 600 sampled trajectories are drawn from the distribution highlighted in Table 3. To emulate target attitude variations in each Monte Carlo case, the true differential ballistic coefficient is modeled as a sinusoid with a mean drawn from $\mathcal{U}(-100\%, 100\%)$, an amplitude of 10%, and a frequency equal to the target mean motion. The filter employed is a numerical UKF variant that models a 5×5 Earth gravity field, lunisolar gravitation and nonconservative perturbations (drag or SRP as functions of ballistic properties). The estimates of ΔB and the




a) Estimation error in $\delta\lambda$ with and without state augmentation

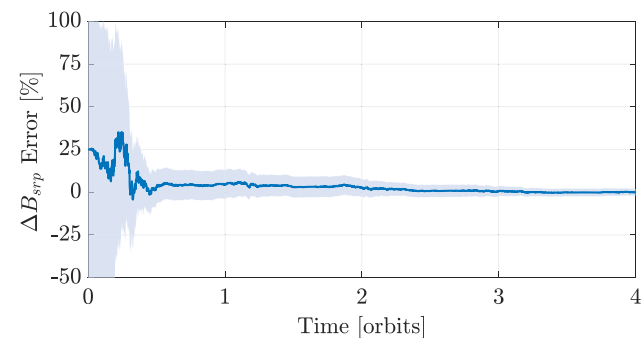


b) Estimation error in drag ΔB with state augmentation

Fig. 11 Relative orbit and differential drag ballistic coefficient mean (solid) and $1\text{-}\sigma$ (shaded) estimation error trends for the *LEOI* Monte Carlo simulations.



a) Estimation error in $\delta\lambda$ with and without state augmentation



b) Estimation error in SRP ΔB with state augmentation

Fig. 12 Relative orbit and differential SRP ballistic coefficient mean (solid) and $1\text{-}\sigma$ (shaded) estimation error trends for the *GEO1* Monte Carlo simulations.

ROEs are initialized with 25% error relative to their true values, and τ is chosen as 10 and 160 min for the LEO and GEO cases, respectively, to effectively sample the parameter over the orbital period. Note that for the *LEO1* case, only ΔB_{drag} is considered; for the *GEO1* case, only ΔB_{srp} is considered. Figures 11a and 12a show the mean and $1\text{-}\sigma$ estimation error trends for the relative mean longitude both including and omitting (by not conducting state augmentation) the ΔB estimation. In the omitting cases, holding the differential ballistic coefficient constant at the erroneous initial conditions yields filter convergence to poor accuracy on the order of 60–65% of the true state value. Instead if the filter is allowed to estimate the differential ballistic coefficient, the estimation of relative mean longitude consistently converges to within approximately $1.4 \pm 4.1\%$ for the *LEO1* case, and $0.7 \pm 4.3\%$ for the *GEO1* case. The estimation error trends for the differential ballistic coefficient state in each of these two cases are shown in Figs. 11b and 12b, respectively. In both cases, the filter is able to estimate the drag or SRP ballistic properties to well within 1% after an initial transient phase. Because the ballistic coefficient provides insight into spacecraft area-to-mass ratio, attitude variation, and material properties, an accurate estimate from the filter provides new information to the observer to be used for meeting mission requirements.

D. Absolute Orbit Determination and Multi-Agent Scenarios

Keeping with the theme of extending filter utility and flexibility, this section describes two new features built into the angles-only navigation framework. In the first, the filtering methodology is extended and assessed for multi-agent scenarios involving several targets and/or observers. In the second feature, state augmentation is used to simultaneously estimate target orbital motion as well as the absolute orbit of the observer. The approach to achieve these goals is useful for formation and swarm sizes that are of practical scientific and engineering impact. In particular, the capability to simultaneously estimate the observer absolute state and the relative state of multiple targets is a key enabling technology and a crucial design trade for swarm missions away from Earth orbit (e.g., about Mars) where absolute orbit information is sparse and inaccurate. Accordingly,

these navigation features are foundational components of the StarFOX experimental navigation payload discussed in Sec. IV.

For the problem of extending the filtering framework to multi-agent scenarios, a pragmatic approach is taken. For the case of multiple targets viewed by a single observer, the estimation state of the filter is simply composed of the ROEs for each of the targets under consideration with respect to a common reference orbit or observer. Incoming measurements consist of pairs of bearing angles for each visible target, with associated measurement correspondences determined based on the performance of the multi-hypothesis target identification and tracking algorithm developed by Kruger and D'Amico [53]. For the multi-observer case where a single target is being estimated, the assumption is that an intersatellite link between the two observers exists for mutual sharing absolute orbit information and observer-specific bearing angle measurements of the target; in this way, the multiple observers estimate the target state in a decentralized manner. Naturally, the general multi-observer and multi-target case is simply a combination of these two approaches where observer formation and decentralized estimation and communication topology can be flexibly reconfigured based on the particular mission application. In this general case, the dynamical system is

$$\begin{aligned} \mathbf{x}(t) &= \left[\delta\mathbf{a}(t)^{(1)}, \dots, \delta\mathbf{a}(t)^{(n_t)} \right]^T \in \mathbf{R}^{6n_t} \\ \dot{\mathbf{x}}(t) &= \left[\mathbf{f}(\delta\mathbf{a}^{(1)}, \mathbf{a}_o, \mathbf{p}, t), \dots, \mathbf{f}(\delta\mathbf{a}^{(n_t)}, \mathbf{a}_o, \mathbf{p}, t) \right]^T \\ \mathbf{y}(t) &= \left[\mathbf{z}_{(1)}^{(1)}, \dots, \mathbf{z}_{(n_o)}^{(1)}, \dots, \mathbf{z}_{(1)}^{(n_t)}, \dots, \mathbf{z}_{(n_o)}^{(n_t)} \right]^T \in \mathbf{R}^{2n_on_t} \end{aligned} \quad (22)$$

where n_t is the number of targets, n_o is the number of observers, and $\mathbf{z}_{(i)}^{(j)}$ is the evaluation of the measurement model \mathbf{h} for the i th observer viewing the j th target:

$$\mathbf{z}_{(i)}^{(j)} = \mathbf{h} \left(\delta\mathbf{a}^{(j)}, \mathbf{a}_{(i)}, \mathbf{v}R_{(i)}^T \right) \quad (23)$$

Although this new feature will be rigorously assessed in the subsequent section, an insightful lesson can be learned from the following scenario. The *LEO2* observer (denoted observer 1, *O1*) is attempting to estimate the ROEs of a target at a separation of 100 km in the anti-link direction, with the true ROEs given by

$$\delta\mathbf{a}_t = (0, -100.0, 1.2, 0, 0.8, 0.8)^T \text{ km}$$

For approximately one orbit, *O1* attempts to estimate this target state while establishing a link with a secondary observer (*O2*) located in a relative formation with the primary observer prescribed by *RO1* in Table 3. Once the communication link is established, *O1* and *O2* exchange their respective measurements of the target, as well as their absolute orbits, and compute an estimate of the targets orbit using the numerical UKF variant receiving measurements every 120 s. The resulting relative mean longitude error as estimated by *O1* is shown in Fig. 13. As soon as measurements are shared, the estimation error quickly drops to within $0.20 \pm 0.35\%$ of the true value and remains stable for the remainder of the coordinated observation phase. This result




Fig. 13 Observer 1 filter $a\delta\lambda$ estimation error and $3\text{-}\sigma$ covariance bounds. Link with observer 2 occurs at red dashed line.

makes intuitive sense when returning to the fundamental observability constraints: by gathering simultaneous measurements from multiple vantage points, the target is uniquely localized with respect to the observers with accuracy approaching the noise floor of the sensors.

Lastly, the filtering methodology is extended to enable simultaneous estimation of target relative motion and the observer absolute orbit. Recall from Eq. (8) that the ROE state is intrinsically evolving as a function of both the absolute orbital state of the observer, as well as due to the differential accelerations acting on the target–observer pair. Rather than treating the observer state as known in this context, a natural extension of the filter is to instead append the absolute OEs of the observer to the estimation state. In this context, relative bearing angle measurements are used to supplement a sparse (or completely nonexistent) absolute orbit update from an external source like GNSS, TLEs, or the Deep Space Network. The dynamical system becomes

$$\begin{aligned} \mathbf{x}(t) &= [\delta\boldsymbol{\alpha}(t), \boldsymbol{\alpha}_o(t)]^T \in \mathbf{R}^{12} \\ \dot{\mathbf{x}}(t) &= \left[\mathbf{f}(\delta\boldsymbol{\alpha}(t), \boldsymbol{\alpha}_o(t), \mathbf{p}(t), t), G(\boldsymbol{\alpha}_o) \mathbf{d}_o^R \right]^T \\ \mathbf{y}(t) &= \left[\mathbf{z}_{(1)}^{(1)}, \boldsymbol{\alpha}_{o, \text{meas}} \right]^T \in \mathbf{R}^8 \end{aligned} \quad (24)$$

Of course, requiring additional tightly coupled states to be estimated necessitates a consideration of the system observability in order to assess whether it is even feasible to do so. The work of Hu et al. [54] demonstrates an observability analysis of two- and three-satellite formations that share inertial bearing angle measurements of one another (note that this is essentially equivalent to solving for the relative orbit of a target and the absolute orbit of an observer). In particular, the assessment is formulated using a Lie derivatives approach with the inertial satellite dynamics parameterized in Cartesian coordinates. A key lesson learned from that work was that the two-satellite formation joint orbit determination is weakly unobservable, whereas the three-satellite formation may enable weakly observable orbit determination under certain configurations. Although this lack of general observability in multi-satellite angles-only absolute orbit determination might appear to pose a limitation for this paper, there are several reasons why this problem is still relevant. A fundamental pillar of this research is that the formulation of the estimation problem in ROE space intrinsically improves the overall observability of the system by better exploiting the connection between the measurement model and the nonlinear curvilinear nature of orbital motion; accordingly, it stands to reason that this same philosophy can improve the complete absolute/relative orbit determination problem. Indeed, the observability constraint decouples into an ambiguity between the observer mean argument of latitude and the relative mean longitude. As a result, in scenarios where good observer absolute orbit information is available or where good relative orbit knowledge is available, the observability limitation is of no practical significance; it is only in cases where both absolute and relative orbit knowledge is poor that it can become an issue. However, because the ROEs are slowly evolving under the effect of perturbations (and in fact constant in strictly Keplerian orbits), any possible estimation divergence due to weak unobservability is expected to occur over very large time spans as compared with the results presented in [54] for the Cartesian states. As such, the simultaneous estimation of target relative orbits and observer absolute orbit may still be practically feasible with very sparse external absolute orbit updates to correct the slow filter drift. Similarly, the use of measurement sharing between multiple coordinated observers can immediately resolve the unobservable ambiguity.

IV. StarFOX: The Starling Formation-Flying Optical Experiment

In light of the huge performance and flexibility improvements that can be obtained in multi-agent scenarios, the navigation framework described in Sec. III will be further assessed in the context of the planned StarFOX experiment, which will be conducted as part of the NASA Starling-1 mission [24]. The Starling-1 mission consists of a swarm of four 6U CubeSats that will be launched into LEO in 2022.

The primary objectives of the mission are to demonstrate new technologies in ad-hoc networking, cluster flight, and optical navigation in space. The latter objective is met by conducting the StarFOX experiment, which builds on the contributions of previous experiments including ARGON [18] and AVANTI [20].

Previous angles-only navigation experiments have been affected by three common limitations: 1) reliance on accurate prior information to initialize the navigation filter (e.g., use of NORAD TLEs for all spacecraft), 2) reliance on frequent translational maneuvers to rectify observability limitations, and 3) use of only two spacecraft. These limitations prevent the use of angles-only techniques for navigation systems in many mission applications of interest (e.g., science missions orbiting other planetary bodies or in deep space). Specifically, reliance on accurate prior information precludes use of angles-only navigation in orbit regimes where accurate reference metrology (e.g., GNSS or TLE) is unavailable. Reliance on maneuvers introduces a complex coupling between the navigation and control systems and reduces mission lifespans. Finally, use of more than two spacecraft requires development of new algorithms that can provide accurate and robust target identification across sequences of images from multiple sensors.

The StarFOX experiment aims to remove all three of these limitations, allowing a swarm of spacecraft to perform absolute and relative navigation using only a single coarse absolute orbit estimate and bearing angles from one or more observers. This degree of generality would enable use of angles-only navigation in a wide range of missions of interest to the scientific community (e.g., performing planetary science in Mars orbit) while adding minimal load to resources like the Deep Space Network.

A. Scenario Definitions and Rationale

StarFOX includes an extensive experiment campaign that will characterize the performance of angles-only navigation algorithms in a wide range of scenarios. These scenarios include different formation geometries, number of observers, frequency and accuracy of absolute orbit estimates, and filter modes (i.e., numerical vs analytical dynamics, process noise handling, etc.). For the purposes of this paper, two test cases have been selected to demonstrate the suitability of the described filter architecture for use in challenging scenarios with multiple targets and only a single absolute orbit estimate.

For these tests, the selected filter configuration is informed by the findings in the previous section. Specifically, the filter estimates a state consisting of its own osculating OEs (i.e., the observer absolute state) and the osculating relative orbits of each target in the field of view. Leveraging this state definition, the dynamics model numerically integrates the GVEs using a Runge–Kutta fourth-order integrator. The onboard dynamics model includes a 20×20 earth gravity field and atmospheric drag based on the Harris–Priester atmospheric density model [55]. The computation cost associated with orbit propagation is minimized by adopting the procedure described in [9], which exploits the triangular structure of the matrix square root to reduce the number of calls to the orbit propagator by a factor of nearly two. The filter uses adaptive process noise tuning to improve convergence speed and robustness. Although the results in the previous section show that parameters such as sensor biases and ballistic coefficients can be estimated, these features are not included in these scenarios because 1) the Starling-1 spacecraft are equipped with star trackers that will be calibrated before launch, and 2) the formation is deployed at an altitude of 500 km, where differential drag between nearly identical spacecraft will be small.

The selected simulations are one day in duration and do not include translational maneuvers for any spacecraft. The simulations use identical formation geometries and differ only in the quality of the initialization of the absolute orbit and use of the intersatellite link to share measurements between multiple observers. The initial osculating absolute orbit of the observer and relative orbits of the three targets are provided in Table 5 where $T1$, $T2$, and $T3$ denote targets 1, 2, and 3, respectively. These orbits were selected to provide a baseline separation of 1 km in the plane perpendicular to the line of sight using relative eccentricity/inclination vector separation [56],

Table 5 Initial osculating orbital elements for HITL simulations

Observer	\tilde{a} (km)	\tilde{e}_x	\tilde{e}_y	\tilde{i} (deg)	$\tilde{\Omega}$ (deg)	\tilde{u} (deg)
	6868	5.1×10^{-4}	4.5×10^{-4}	98.0	0.0	0.0
Relative orbit (RO)	$a\tilde{\delta}a$ (km)	$a\tilde{\delta}l$ (km)	$a\tilde{\delta}e_x$ (km)	$a\tilde{\delta}e_y$ (km)	$a\tilde{\delta}i_x$ (km)	$a\tilde{\delta}i_y$ (km)
T1: Target 1	0	65.0	0	1.0	0	1.0
T2: Target 2	0	133.0	1.2	1.6	1.2	1.6
T3: Target 3	0	200.0	0	-1.0	0	-1.0

thereby minimizing the number of maneuvers required for station-keeping. The distinct features of each test case are described in the following.

1. Test Case 1: Single Observer

In the first simulation, a single observer estimates its own absolute orbit and the relative orbits of three targets. Although it was shown by Hu et al. that a single observer cannot completely resolve the state [54], this scenario is still relevant if divergence of the state estimate is slow. Specifically, the unobservable mode(s) can be corrected by sparse updates of the observer's absolute orbit (e.g., using the Deep Space Network [57]). As such, the goal of this simulation is to show that a single estimate of the absolute orbit of the observer is sufficient to enable estimation of the complete state of the swarm with useful accuracy for one day.

To demonstrate this capability, the filter is initialized with an estimate of the absolute orbit of the observer with accuracy consistent with GNSS or Deep Space Network [57] solutions (10 m, 1 cm/s $1-\sigma$ errors in each position and velocity component, respectively). The provided relative orbit estimates are subject to errors and $1-\sigma$ uncertainties of 10 km in δl and 100 m in all other ROEs, which are representative of the expected accuracy of current batch relative orbit determination algorithms [22]. The navigation algorithms are provided with measurements from the onboard camera every 120 s for 70% of each orbit. No measurement updates are performed in the remaining 30% of each orbit to simulate outages due to eclipse, blinding of the sensor, or mission operations constraints.

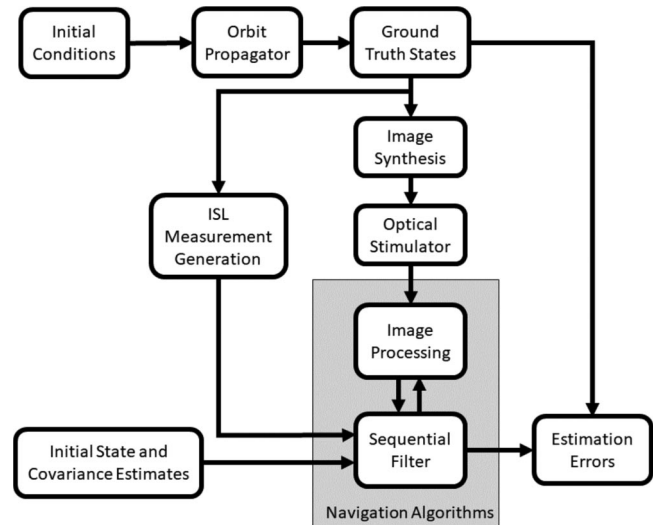
2. Test Case 2: Multiple Observers

In the second test case, two observers share bearing angle measurements to estimate the complete state of the swarm from a coarse initialization of the absolute and relative orbits. The second observer is T3, which has a camera pointed in the anti-flight direction (allowing it to see the primary observer, T1, and T2). The objectives of this test are twofold. First, it will show that the weak local observability shown by Hu et al. [54] for Keplerian orbits is sufficient to enable estimation of the complete state of a swarm in the presence of representative perturbations and errors. Second, it will demonstrate that angles-only navigation algorithms provide sufficient navigation accuracy to enable autonomous swarm operations around other planetary bodies with minimal ground interaction.

To demonstrate these capabilities, the filter is provided with an absolute orbit estimate and covariance that are consistent with TLE errors (1 km error in semimajor axis, and several km of error in other elements). As in test case 1, the provided relative orbit estimates are subject to errors and $1-\sigma$ uncertainties of 10 km in δl and 100 m in all other ROEs. The navigation algorithms are provided with measurements from both observers every 120 s for 70% of the orbit and no measurements are provided in the remaining 30% of the orbit.


B. Simulation Environment

To ensure that errors in inputs for StarFOX are as representative as possible of errors encountered in flight, the previously described simulation architecture is augmented to include a Nano Star Tracker from Blue Canyon Technologies [58] in the loop. This is the same sensor that will be used on the Starling-1 satellites. The augmented simulation environment is illustrated in Fig. 14 and described in the following.

**Fig. 14** Illustration of HITL simulation environment.

First, the initial conditions (initial epoch, sample time, final epoch, initial osculating orbits, and sensor attitude profile for each spacecraft) are specified by the user. Next, the ground truth states (orbit and attitude) at each epoch are computed by propagating the orbits of all four spacecraft using the high-fidelity numerical orbit propagator described in Sec. II, which includes all significant perturbations in LEO.

The ground truth data are used to generate real images of the target spacecraft and background stellar/nonstellar objects at each measurement epoch using the Space Rendezvous Laboratory's optical stimulator testbed [59]. This is accomplished in two steps. First, the optical stimulator generates high-fidelity synthetic images of the target spacecraft and background stars at each measurement epoch using 3D vector graphics in OpenGL. Stellar objects and far-range targets are represented by Gaussian point spread functions (PSF). Five random PSF are also added to each image to simulate unidentified objects in the field of view, positioned according to a uniform distribution. These synthetic images are converted to real images using the hardware component of the optical stimulator, which is shown in Fig. 15. The optical stimulator hardware includes the sensor under test, two lenses, and an OLED microdisplay with a high dynamic-range, which stimulates the sensor by rendering synthetic images of the space environment. The lenses and display are movable to enable variable magnification

**Fig. 15** CAD model of the optical stimulator.

and the system is calibrated to ensure that the geometry and radiosity of simulated imagery closely matches that of the space environment. Specifically, angular residuals between desired and measured feature locations are less than 10 arcsec. Design, calibration, and use of the optical stimulator are fully described in [59].

Next, the image processing algorithm takes the time-tagged images from the star tracker and computes the sensor attitude and bearing angles to each target at each epoch. This is accomplished by a simple four-step procedure. First, illuminated pixel clusters are centroided and converted to unit vectors in the sensor frame using a calibrated camera model. Second, stars in the image are identified using the Pyramid algorithm [60]. Third, the identified stars and corresponding pointing vectors are used to compute the sensor attitude. Finally, bearing angle measurements to each target are assigned using the Spacecraft Angles-only MUlti-target-tracking Software (SAMUS) algorithm. SAMUS applies principles of multi-hypothesis tracking, using a combination of state information provided by the navigation filter and the kinematics of target trajectories (i.e., the geometric features of Figs. 1a and 1b) observed across previous images to robustly distinguish and assign measurements to multiple targets. The details of SAMUS and the image processing algorithm are described in [53].

Finally, the sequential filter uses the provided initial state estimate and covariance and the bearing angle measurements from image

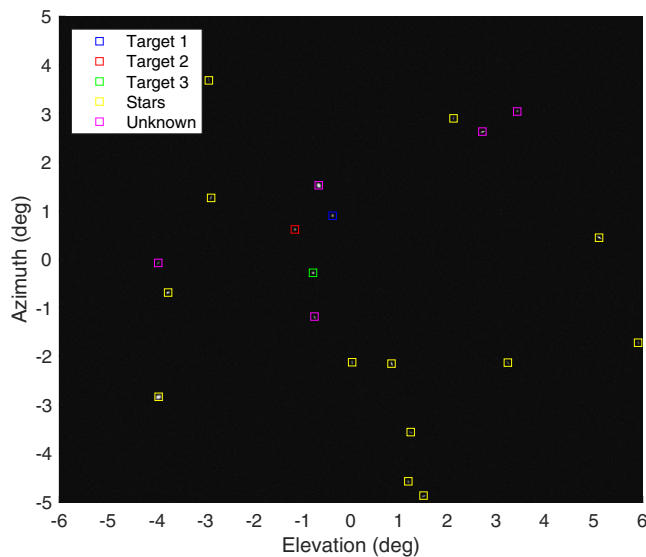


Fig. 16 Test case 1 initial input image and object classifications after processing by SAMUS.

processing (and the intersatellite link if provided) to compute estimates of the absolute and relative orbits of the swarm and associated uncertainties at each epoch. It is assumed that the measurement correspondence problem for measurements received over the intersatellite link (i.e., determination of which bearing angle measurement from the intersatellite link corresponds to each target seen by the observer) is solved. Work is underway to generalize these algorithms to solve this problem for the flight experiment.

C. Results

1. Image Processing

The performance of the image processing algorithm in regard to multi-target measurement assignment is described using the traditional metrics of precision, recall, and accuracy. Precision quantifies assignment robustness and the number of false-positive assignments; recall quantifies assignment frequency and the number of false-negative assignments; and accuracy treats overall assignment performance. Figure 16 presents an example output of image processing, displaying the classifications of detected objects. Unknown objects are those added by the optical stimulator to emulate unidentified stellar or nonstellar objects.

In test case 1, SAMUS produces a precision of 99.93%, recall of 97.25%, and accuracy of 99.45% after classifying a total of 10,375 centroided objects. The single false positive results in a maximum bearing angle assignment error of 98 arcsec. This error is caused by the temporary overlap of a target pixel cluster with a star pixel cluster, leading to an inaccurate centroid location. In test case 2, SAMUS produces a precision of 99.87%, recall of 97.25%, and accuracy of 99.44% over 10,375 centroided objects. An additional false positive is observed at the beginning of the simulation due to large initial uncertainties in the estimated state, which creates similarly large uncertainties in the predicted bearing angle measurements. Multiple bearing angles can then be reasonable candidates for a target measurement and may result in an incorrect assignment. In the case of large initial covariances, SAMUS may instead choose to track targets using kinematic means, without using filter state information. This causes a maximum bearing angle error of 129 arcsec with two false positives in total. Overall, however, image processing produces highly accurate outputs for multi-target measurement assignment. The correct bearing angles are passed to the filter, and consequently, the filter state effectively aids assignment in new images.

2. Test Case 1: Single Observer

The estimation errors and uncertainties for the relative orbits in test case 1 are shown in Fig. 17. It can be seen that the filter converges to errors of less than 200 m in $a\delta\lambda$ ($<0.4\%$) and less than 10 m in all other ROEs within a few orbits. The steady-state range uncertainty for all

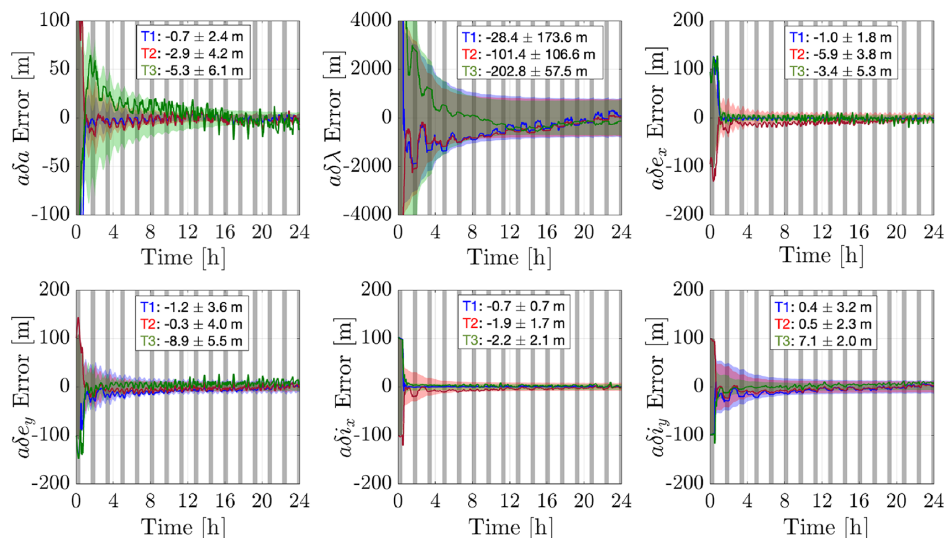


Fig. 17 Relative orbit estimation error and formal covariance for T1, T2, and T3 in test case 1 with measurement blackout periods indicated in gray.

targets is 1 km (<1.5%). These results are consistent with the filter performance shown in Sec. III.

With this in mind, consider the absolute orbit estimation error and uncertainty shown in Fig. 18. This figure includes results for simulations with (blue) and without (red) measurement updates applied to the absolute orbit to provide a quantitative metric of the efficacy of using angles measurements for absolute orbit estimation with a single observer. The errors are plotted as ROEs of the estimated state with respect to the ground truth state to provide a simple geometric interpretation of their meaning. That is, $a\delta a$ represents the error in semimajor axis, $a\delta\lambda$ captures the error in the mean argument of latitude, $a\delta e_x$ and $a\delta e_y$ capture errors in the eccentricity vector estimate, $a\delta i_x$ represents errors in the inclination estimate, and $a\delta i_y$ represents errors in the RAAN estimate. The corresponding uncertainties in ROE space are computed by applying an unscented transform to the absolute orbit covariance matrix.

It is clear from Fig. 18 that without measurement updates applied to the absolute orbit, the uncertainty in the orbit semimajor axis grows by several hundred meters over one day. As a result of the coupling between orbital energy and orbital phasing, this divergence in semimajor axis leads to rapidly growing uncertainty and errors in the mean argument of latitude. On the other hand, when measurement updates are applied to the absolute orbit, the uncertainty in the semimajor axis is bounded to about 100 m, which is consistent with the behavior of the error. This result, in turn, allows the error in the mean argument of latitude to be bounded to a few hundred meters. The measurement updates also reduce the growth of errors in the eccentricity and inclination vector by approximately 50%. Although these results do indicate marginal convergence in the estimate of the absolute orbit (as expected from results in [54]), it is noteworthy that the degradation is very slow (<1 km/day) and this accuracy is sufficient to achieve relative navigation accuracy of 200 m or better.

The steady-state estimation errors expressed in relative position and velocity in the local RTN frame for this test case are provided in Table 6. These results show that the absolute position and velocity are estimated to within 1 km and 1 m/s in position and velocity. The relative positions along the radial (\mathcal{R}_x) and cross-track (\mathcal{R}_z) axes are estimated to within 25 m or better and the relative velocity components are estimated to within 20 mm/s. The along-track (\mathcal{R}_y) component of the relative positions (which is the dominant component of the separation) is estimated to within 1 km.

Finally, the pre- and post-fit bearing angle residuals for each target in this simulation are shown in Fig. 19. It is clear from this plot that the measurement residuals for all targets converge to 25 arcsec or smaller within three orbits, in agreement with the convergence trends observed in Figs. 17 and 18.

Overall, these results suggest that a single observer can estimate the complete state of a swarm without translational maneuvers if it is

Table 6 Test case 1 steady-state estimation errors ($1\text{-}\sigma$) in RTN

Component	Observer	Target 1	Target 2	Target 3
$\delta r_x^{\mathcal{R}}$ (m)	26 ± 437	1.0 ± 9.6	1.5 ± 15.7	4.3 ± 23.2
$\delta r_y^{\mathcal{R}}$ (m)	69 ± 902	-27 ± 801	-102 ± 759	-205 ± 722
$\delta r_z^{\mathcal{R}}$ (m)	224 ± 651	0.6 ± 8.7	0.0 ± 9.2	-0.2 ± 6.3
$\delta v_x^{\mathcal{R}}$ (mm/s)	112 ± 451	0.4 ± 8.9	-0.7 ± 9.3	-1.8 ± 7.8
$\delta v_y^{\mathcal{R}}$ (mm/s)	-68 ± 962	-1.2 ± 18.0	1.5 ± 20.0	3.5 ± 18.1
$\delta v_z^{\mathcal{R}}$ (mm/s)	102 ± 527	-1.0 ± 9.8	-0.7 ± 10.6	-0.5 ± 6.9

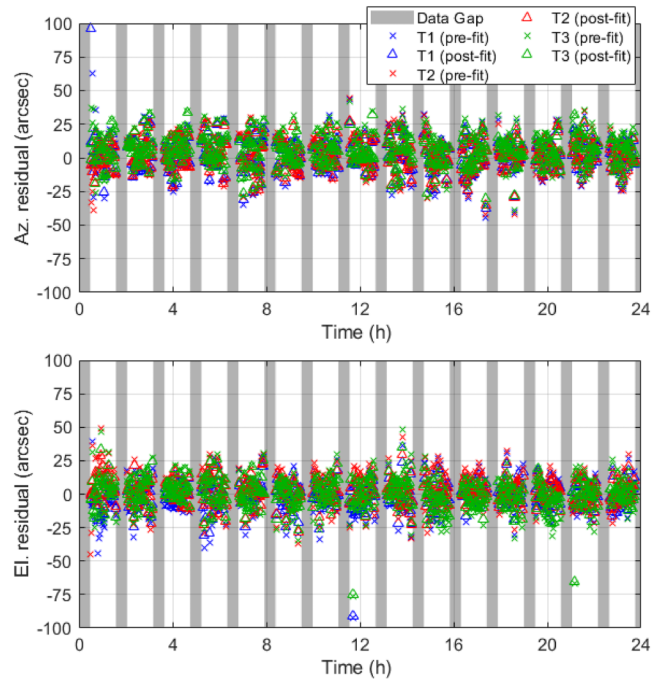


Fig. 19 Pre-fit and post-fit bearing angle residuals for each target in test case 1.

provided with infrequent absolute orbit estimates (e.g., from the Deep Space Network).

3. Test Case 2: Multiple Observers

The estimation error and uncertainty for the absolute orbit of the observer are shown in Fig. 20. As before, these errors are cast in ROE

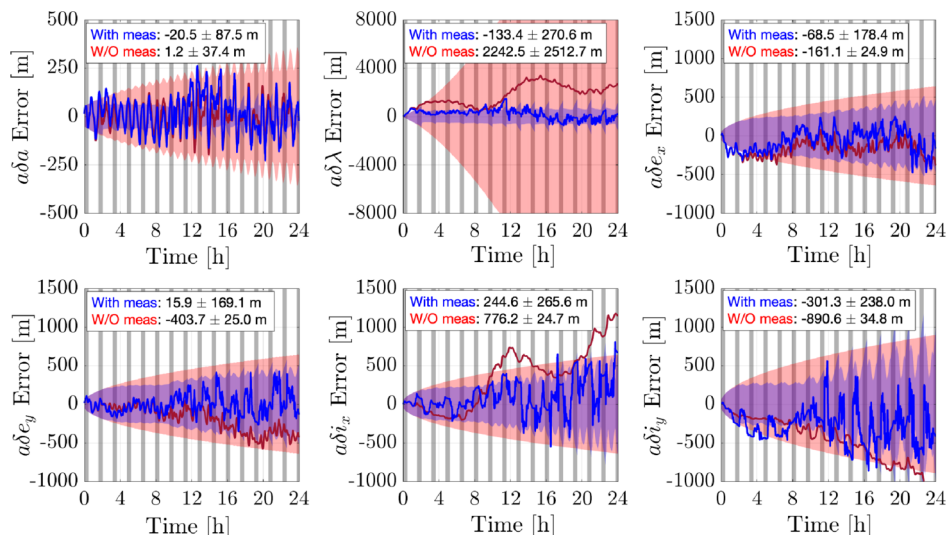


Fig. 18 Absolute orbit propagation error (red) and estimation error (blue) with corresponding uncertainties. Measurement blackout periods are shown in gray.

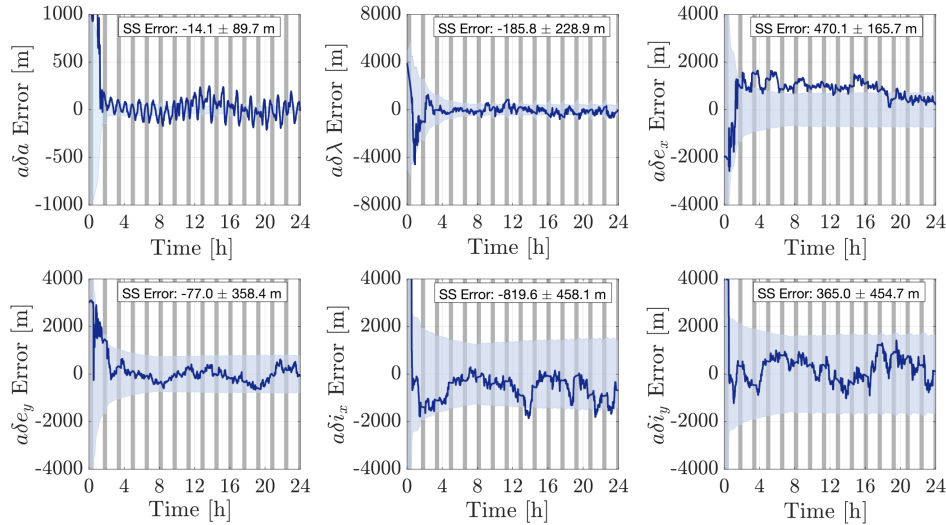


Fig. 20 Absolute orbit estimation error and formal covariance with measurement blackout periods shown in gray.

space to provide a simple geometric interpretation. With the added information from $T3$'s bearing angle measurements, the filter is able to converge to steady state in only three orbits with errors of less than 100 m in the semimajor axis, less than 500 m in the mean argument of latitude and eccentricity vector, and about 1 km in the inclination and RAAN. These larger errors in inclination and RAAN are expected as they correspond to a 30 arcsec rotation of the orbit plane; the effect of this rotation on the observed measurements is comparable to the measurement noise.

The estimation errors and uncertainties of the relative orbits of the targets are shown in Fig. 21. As in the previous simulation, the filter is able to converge to steady-state errors of less than 200 m ($< 0.4\%$) in $a\Delta \lambda$ and less than 10 m in $a\Delta a$, $a\Delta e_x$, and $a\Delta e_y$. However, the estimates of the relative inclination vectors are slightly degraded, with steady-state errors of up to 25 m. This is likely due to the increased uncertainty in the inclination and RAAN, which is twice as large as the final uncertainty from Test Case 1. Finally, it can be seen that the uncertainties of the relative state estimates exhibit a clear linear scaling with separation (because $T1$ is the closest and $T3$ is the furthest from the observer). This range dependence is expected because position estimation error scales linearly with separation when using angles measurements.

Overall, these simulation results show that the complete state of a swarm (absolute and relative orbits) can be accurately estimated from a single coarse initial estimate that is refined using bearing angle measurements from two observers. This approach does not require

translational maneuvers from any spacecraft, though inclusion of known maneuvers would improve convergence speed and estimation accuracy. Additionally, the results of both test cases suggest that the StarFOX experiment will be able to provide a first-of-a-kind demonstration of angles only navigation for spacecraft swarms that does not require maneuvers or relative state estimates from the ground.

V. Conclusions

This research has put forth a complete maneuver-free angles-only navigation approach for deployment in multi-satellite systems operating in planetary orbit regimes. By formulating the estimation architecture using an ROE state description, and using key nonlinearities in the unscented Kalman filtering framework, excellent performance, robustness, and flexibility were shown in several Monte Carlo simulations. There were several key lessons learned about the capabilities of the navigation algorithms. As a broad theme, it was shown that filter flexibility is enabled by expanding the estimation state. In particular, sensor biases can be estimated as first-order Gauss–Markov processes, which allows on-orbit calibration of the crucial navigation sensor to be achieved online, and enables improved robustness and measurement fault detection. Additionally, target satellite ballistic properties, like the drag or SRP ballistic coefficients, were shown to be obtainable from angles-only measurements with an appropriate filter dynamics model choice. These ballistic parameters provide a wealth of additional information about the target structural properties and attitude

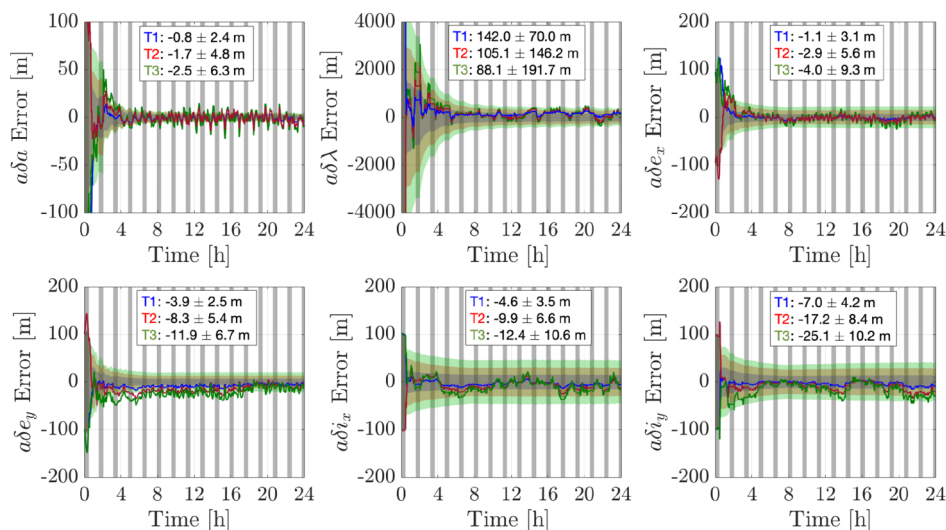


Fig. 21 Relative orbit estimation error and formal covariance for $T1$, $T2$, and $T3$ in test case 2 with measurement blackout periods indicated in gray.

variation, and sequentially improve the dynamics modeling. Furthermore, multi-target angles-only estimation was made feasible by pragmatic expansion of the single-target filtering methodology to multiple measurement inputs with known correspondences. This allowed a single observer satellite to scan a large vicinity and determine the orbital states of many targets, enabling new capabilities in space situational awareness and formation flying. In a similar multi-satellite vein, coordinated use of multiple observers employing the navigation algorithms in a decentralized fashion allowed extremely rapid and accurate estimation convergence while maintaining flexibility in formation and communication topology design to meet additional mission parameters. Lastly, the use of relative angles-only measurements even enabled the simultaneous estimation of the absolute orbit of an observer spacecraft. Although the observability problem is even more limiting in this case, the use of ROEs in the framework indicates that potential filter instabilities occur over much larger time horizons than traditional Cartesian-coordinate-based filters. Accordingly, angles-only measurements can be used to supplement a very sparse and inaccurate absolute orbit update for the observer in order to achieve practically useful absolute orbit estimation performance. In light of these newly verified algorithmic advantages, a partnership with NASA Ames Research Center was formed to develop the novel on-orbit angles-only navigation testbed known as the StarFOX. The core intent of this experiment is to increase the technology readiness level of the proposed angles-only navigation prototype for future interplanetary multi-satellite formations.

Several current research efforts are underway that build from the work presented in this paper. In particular, robust methods for handling the multi-target measurement identification and correspondence are being formulated to complete the full angles-only navigation system. Similarly, the problem of simultaneous absolute and relative orbit determination is being examined both from a theoretical observability standpoint as well as in rigorous Monte Carlo evaluation in order to understand the practical capabilities of this concept for full swarm localization outside of Earth orbit. Lastly, the hardware-in-the-loop verification achieved in this work is being expanded to include a dedicated porting and embedding effort to run the algorithms on representative spacecraft processors in preparation for the on-orbit StarFOX experiment commissioning in 2022.

Acknowledgments

This work was supported by the Air Force Research Laboratory's Control, Navigation, and Guidance for Autonomous Spacecraft (CoNGAS) contract FA9453-16-C-0029, and the NASA Small Spacecraft Technology Program (cooperative agreement number 80NSSC18M0058) for contribution to the Starling-1/StarFOX efforts. Additionally, the authors would like to acknowledge the work of Eric Stoker-Spirt on the research and development of algorithms for absolute navigation.

References

- [1] Tapley, B. D., Bettadpur, S., Watkins, M., and Reigber, C., "The Gravity Recovery and Climate Experiment: Mission Overview and Early Results," *Geophysical Research Letters*, Vol. 31, No. 9, 2004. <https://doi.org/10.1029/2004GL019920>
- [2] Krieger, G., Moreira, A., Fiedler, H., Hajsek, I., Werner, M., Younis, M., and Zink, M., "TanDEM-X: A Satellite Formation for High-Resolution SAR Interferometry," *IEEE Transactions on Geoscience and Remote Sensing*, Vol. 45, No. 11, 2007, pp. 3317–3341. <https://doi.org/10.1109/TGRS.2007.900693>
- [3] D'Amico, S., Ardaens, J.-S., and Larsson, R., "Spaceborne Autonomous Formation-Flying Experiment on the PRISMA Mission," *Journal of Guidance, Control, and Dynamics*, Vol. 35, May 2012, pp. 834–850.
- [4] Burch, J., Moore, T., Torbert, R., and Giles, B., "Magnetospheric Multi-scale Overview and Science Objectives," *Space Science Reviews*, Vol. 199, Nos. 1–4, 2016, pp. 5–21. <https://doi.org/10.1007/s11214-015-0164-9>
- [5] Chung, S.-J., and Hadaegh, F. Y., "Swarms of Femtosats for Synthetic Aperture Applications," *Proceedings of the 4th International Conference on Spacecraft Formation Flying Missions & Technologies, Advances in the Astronautical Sciences*, National Research Council Canada, Ottawa, May 2011.
- [6] Morgan, D., Chung, S.-J., Blackmore, L., Acikmese, B., Bayard, D., and Hadaegh, F. Y., "Swarm-Keeping Strategies for Spacecraft Under J_2 and Atmospheric Drag Perturbations," *Journal of Guidance, Control, and Dynamics*, Vol. 35, No. 5, 2012, pp. 1492–1506. <https://doi.org/10.2514/1.55705>
- [7] Morgan, D., "Guidance and Control of Swarms of Spacecraft," Ph.D. Thesis, Univ. of Illinois Urbana-Champaign, Champaign, IL, 2015.
- [8] D'Amico, S., "Autonomous Nanosatellite Swarming Using Radio-Frequency and Optical Navigation (ANS)," NASA Fact Sheet, Stanford Space Rendezvous Lab (SLAB), 2018, https://damicos.people.stanford.edu/sites/g/files/sbiybj2226/ff/sstp_fact_sheet_ans_2019.pdf.
- [9] Stacey, N., and D'Amico, S., "Autonomous Swarming for Simultaneous Navigation and Asteroid Characterization," *AAS/AIAA Astrodynamics Specialist Conference*, AAS Paper 18-448, San Diego, CA, 2018.
- [10] Lippe, C., and D'Amico, S., "Spacecraft Swarm Dynamics and Control About Asteroids," *2019 International Workshop on Satellite Constellations and Formation Flying*, 2019.
- [11] Woffinden, D. C., and Geller, D. K., "Observability Criteria for Angles-Only Navigation," *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 45, No. 3, 2009, pp. 1194–1208. <https://doi.org/10.1109/TAES.2009.5259193>
- [12] Gaias, G., D'Amico, S., and Ardaens, J.-S., "Angles-Only Navigation to a Noncooperative Satellite Using Relative Orbital Elements," *Journal of Guidance, Control, and Dynamics*, Vol. 37, March 2014, pp. 439–451. <https://doi.org/10.2514/1.61494>
- [13] Sullivan, J., Koenig, A., and D'Amico, S., "Improved Maneuver-Free Approach to Angles-Only Navigation for Space Rendezvous," *26th AAS/AIAA Space Flight Mechanics Meeting*, AAS Paper 16-530, San Diego, CA, 2016.
- [14] Woffinden, D. C., and Geller, D. K., "Optimal Orbital Rendezvous Maneuvering for Angles-Only Navigation," *Journal of Guidance, Control, and Dynamics*, Vol. 32, July 2009, pp. 1382–1387. <https://doi.org/10.2514/1.45006>
- [15] Grzymisch, J., and Fichter, W., "Observability Criteria and Unobservable Maneuvers for In-Orbit Bearings-Only Navigation," *Journal of Guidance, Control, and Dynamics*, Vol. 37, July 2014, pp. 1250–1259. <https://doi.org/10.2514/1.62476>
- [16] Franquiz, F. J., Muñoz, J. D., Udrea, B., and Balas, M. J., "Optimal Range Observability Maneuvers of a Spacecraft Formation Using Angles-Only Navigation," *Acta Astronautica*, Vol. 153, Feb. 2018, pp. 337–348.
- [17] Wang, J., Butcher, E. A., and Tansel, Y., "Space-Based Relative Orbit Estimation Using Information Sharing and the Consensus Kalman Filter," *Journal of Guidance, Control, and Dynamics*, Vol. 42, No. 3, 2019, pp. 491–507. <https://doi.org/10.2514/1.G003503>
- [18] D'Amico, S., Ardaens, J.-S., Gaias, G., Benninghoff, H., Schleppe, B., and Jørgensen, J. L., "Noncooperative Rendezvous Using Angles-Only Optical Navigation: System Design and Flight Results," *Journal of Guidance, Control, and Dynamics*, Vol. 36, Nov. 2013, pp. 1576–1595. <https://doi.org/10.2514/1.59236>
- [19] Gaias, G., Ardaens, J.-S., and D'Amico, S., "The Autonomous Vision Approach Navigation and Target Identification (AVANTI) Experiment: Objectives and Design," *9th International ESA Conference on Guidance, Navigation & Control Systems*, 2014.
- [20] Gaias, G., and Ardaens, J.-S., "Flight Demonstration of Autonomous Noncooperative Rendezvous in Low Earth Orbit," *Journal of Guidance, Control, and Dynamics*, Vol. 41, No. 6, 2018, pp. 1337–1354. <https://doi.org/10.2514/1.G003239>
- [21] Ardaens, J.-S., and Gaias, G., "Angles-Only Relative Orbit Determination in Low Earth Orbit," *Advances in Space Research*, Vol. 61, No. 11, 2018, pp. 2740–2760. <https://doi.org/10.1016/j.asr.2018.03.016>
- [22] Sullivan, J., and D'Amico, S., "Nonlinear Kalman Filtering for Improved Angles-Only Navigation Using Relative Orbital Elements," *Journal of Guidance, Control, and Dynamics*, Vol. 40, No. 9, July 2017, pp. 2183–2200.
- [23] Koenig, A. W., Guffanti, T., and D'Amico, S., "New State Transition Matrices for Relative Motion of Spacecraft Formations in Perturbed Orbits," *Journal of Guidance, Control, and Dynamics*, Vol. 40, No. 7, 2017, pp. 1749–1768.
- [24] Sanchez, H., McIntosh, D., Cannon, H., Pires, C., Field, M., Sullivan, J., D'Amico, S., Mall, L., and O'Connor, B., "Starling-1: Swarm Technology Demonstration," *32nd Annual Small Satellite Conference*, Aug. 2018.
- [25] Sullivan, J., Grimberg, S., and D'Amico, S., "Comprehensive Survey and Assessment of Spacecraft Relative Motion Dynamics Models," *Journal of Guidance, Control, and Dynamics*, Vol. 40, No. 8, April 2017, pp. 1837–1859.
- [26] Casotto, S., "The Equations of Relative Motion in the Orbital Reference Frame," *Celestial Mechanics and Dynamical Astronomy*, Vol. 124,

- No. 3, 2016, pp. 215–234.
<https://doi.org/10.1007/s10569-015-9660-1>
- [27] Lawden, D., *Optimal Trajectories for Space Navigation*, Butterworths, London, 1963, p. 75.
- [28] Tschauer, J., and Hempel, P., “Optimale Beschleunigungsprogramme für das Rendezvous-Manöver,” *Acta Astronautica*, Vol. 10, Nos. 5–6, 1964, pp. 296–307.
- [29] Yamanaka, K., and Ankersen, F., “New State Transition Matrix for Relative Motion on an Arbitrary Elliptical Orbit,” *Journal of Guidance, Control, and Dynamics*, Vol. 25, Jan. 2002, pp. 60–66.
<https://doi.org/10.2514/2.4875>
- [30] Branco, J., Olmos, D. E., Castellani, L. T., and Cropp, A., “The Formation Flying Navigation System for Proba 3,” *Aerospace Robotics II*, Springer International Publishing, Cham, Switzerland, 2015, pp. 37–47.
- [31] Clohessy, W. H., and Wiltshire, R. S., “Terminal Guidance System for Satellite Rendezvous,” *Journal of the Aerospace Sciences*, Vol. 27, Sept. 1960, pp. 653–658.
<https://doi.org/10.2514/8.8704>
- [32] Schaub, H., “Relative Orbit Geometry Through Classical Orbit Element Differences,” *Journal of Guidance, Control, and Dynamics*, Vol. 27, Sept. 2004, pp. 839–848.
<https://doi.org/10.2514/1.12595>
- [33] D’Amico, S., “Autonomous Formation Flying in Low Earth Orbit,” Ph.D. Thesis, Delft Univ., Delft, The Netherlands, 2010.
- [34] Alfried, K. T. (ed.), *Spacecraft Formation Flying: Dynamics, Control, and Navigation*, Elsevier Astrodynamics Series, Butterworth-Heinemann/Elsevier, Oxford, 2010.
- [35] Guffanti, T., D’Amico, S., and Lavagna, M., “Long-Term Analytical Propagation of Satellite Relative Motion in Perturbed Orbits,” *27th AAS/AIAA Space Flight Mechanics Meeting*, AAS Paper 17-355, San Diego, CA, 2017.
- [36] Guffanti, T., and D’Amico, S., “Linear Models for Spacecraft Relative Motion Perturbed by Solar Radiation Pressure,” *Journal of Guidance, Control, and Dynamics*, Vol. 42, No. 9, 2019, pp. 1962–1981.
- [37] D’Amico, S., and Montenbruck, O., “Proximity Operations of Formation-Flying Spacecraft Using an Eccentricity/Inclination Vector Separation,” *Journal of Guidance, Control, and Dynamics*, Vol. 29, May 2006, pp. 554–563.
- [38] Koenig, A. W., and D’Amico, S., “Robust and Safe n-Spacecraft Swarming in Perturbed Near-Circular Orbits,” *Journal of Guidance, Control, and Dynamics*, Vol. 41, No. 8, 2018, pp. 1643–1662.
<https://doi.org/10.2514/1.G003249>
- [39] Schaub, H., and Junkins, J., *Analytical Mechanics of Space Systems*, AIAA, Reston, VA, 2009, Chap. 9.
- [40] Gaias, G., Lara, M., and Colombo, C., “Accurate Osculating/Mean Orbital Elements Conversions for Spaceborne Formation Flying,” *27th International Symposium on Space Flight Dynamics (ISSFD)*, Engineers Australia, Melbourne, Australia, 2019, p. 1171.
- [41] Montenbruck, O., and Gill, E., *Satellite Orbits: Models, Methods and Applications*, Springer-Verlag, Berlin, 2012, Chap. 3.
- [42] Picone, J., Hedin, A., Drob, D. P., and Aikin, A., “NRLMSISE-00 Empirical Model of the Atmosphere: Statistical Comparisons and Scientific Issues,” *Journal of Geophysical Research: Space Physics*, Vol. 107, No. A12, 2002, pp. SIA–15.
<https://doi.org/10.1029/2002JA009430>
- [43] Beierle, C., Sullivan, J., and D’Amico, S., “Design and Utilization of the Stanford Vision-Based Navigation Testbed for Spacecraft Rendezvous,” *9th International Workshop on Satellite Constellations and Formation Flying*, Univ. of Colorado Boulder, Boulder, CO, June 2017, pp. 1–4.
- [44] Beierle, C., Sullivan, J., and D’Amico, S., “High-Fidelity Verification of Vision-Based Sensors for Inertial and Far-Range Spaceborne Navigation,” *26th International Symposium on Space Flight Dynamics*, June 2017.
- [45] Julier, S. J., and Uhlmann, J. K., “New Extension of the Kalman Filter to Nonlinear Systems,” *AeroSense ’97, Signal Processing, Sensor Fusion, and Target Recognition VI*, Vol. 3068, edited by I. Kadar, International Soc. for Optics and Photonics, 1997, pp. 182–193.
<https://doi.org/10.1117/12.280797>
- [46] Thrun, S., Burgard, W., and Fox, D., *Probabilistic Robotics*, MIT Press, Cambridge, MA, 2005, pp. 67–71, Chap. 3.
- [47] Myers, K., and Tapley, B., “Dynamical Model Compensation for Near-Earth Satellite Orbit Determination,” *AIAA Journal*, Vol. 13, March 1975, pp. 343–349.
<https://doi.org/10.2514/3.49702>
- [48] Myers, K., and Tapley, B., “Adaptive Sequential Estimation with Unknown Noise Statistics,” *IEEE Transactions on Automatic Control*, Vol. 21, No. 4, 1976, pp. 520–523.
<https://doi.org/10.1109/TAC.1976.1101260>
- [49] Fraser, C. T., “Adaptive Extended Kalman Filtering Strategies for Autonomous Relative Navigation of Formation Flying Spacecraft,” Ph.D. Thesis, Carleton Univ., Ottawa, 2019.
- [50] Boyd, S., and Xiao, L., “Least-Squares Covariance Matrix Adjustment,” *SIAM Journal on Matrix Analysis and Applications*, Vol. 27, No. 2, 2005, pp. 532–546.
<https://doi.org/10.1137/040609902>
- [51] Higham, N. J., “Computing a Nearest Symmetric Positive Semidefinite Matrix,” *Linear Algebra and Its Applications*, Vol. 103, May 1988, pp. 103–118.
[https://doi.org/10.1016/0024-3795\(88\)90223-6](https://doi.org/10.1016/0024-3795(88)90223-6)
- [52] Gillespie, D. T., *Markov Processes: An Introduction for Physical Scientists*, Elsevier, New York, 1991.
- [53] Kruger, J., and D’Amico, S., “Autonomous Angles-Only Multitarget Tracking for Spacecraft Swarms,” *2020 AAS/AIAA Astrodynamics Specialist Conference*, AAS Paper 20-449, San Diego, CA, 2020.
- [54] Hu, Y., Sharf, I., and Chen, L., “Three-Spacecraft Autonomous Orbit Determination and Observability Analysis with Inertial Angles-Only Measurements,” *Acta Astronautica*, Vol. 170, May 2020, pp. 106–121.
<https://doi.org/10.1016/j.actaastro.2020.01.005>
- [55] Harris, I., and Priester, W., “Theoretical Models for the Solar-Cycle Variation of the Upper Atmosphere,” *Journal of Geophysical Research*, Vol. 67, No. 12, 1962, pp. 4585–4591.
<https://doi.org/10.1029/JZ067i012p04585>
- [56] D’Amico, S., and Montenbruck, O., “Proximity Operations of Formation-Flying Spacecraft Using an Eccentricity/Inclination Vector Separation,” *Journal of Guidance, Control, and Dynamics*, Vol. 29, May 2006, pp. 554–563.
- [57] Curkendall, D. W., and Border, J. S., “Delta-DOR: The One-Nanoradian Navigation Measurement System of the Deep Space Network—History, Architecture, and Componentry,” *Interplanetary Network Progress Report*, Vol. 42, May 2013, p. 193.
- [58] Palo, S., Stafford, G., and Hoskins, A., “An Agile Multi-Use Nano Star Camera for Constellation Applications,” *Proceedings of the AIAA/USU Conference on Small Satellites*, Aug. 2013.
- [59] Beierle, C., and D’Amico, S., “Variable Magnification Optical Stimulator for Training and Validation of Spaceborne Vision-Based Navigation,” *Journal of Spacecraft and Rockets*, Vol. 56, No. 4, 2019, pp. 1060–1072.
<https://doi.org/10.2514/1.A34337>
- [60] Mortari, D., Samaan, M. A., Bruccoleri, C., and Junkins, J. L., “The Pyramid Star Identification Technique,” *Navigation*, Vol. 51, No. 3, 2004, pp. 171–183.
<https://doi.org/10.1002/j.2161-4296.2004.tb00349.x>