JOURNAL OF GUIDANCE, CONTROL, AND DYNAMICS

Autonomous Navigation of a Satellite Swarm Using Inter-Satellite Bearing Angles

Justin Kruger* and Simone D'Amico†
Stanford University, Stanford, California 94305

https://doi.org/10.2514/1.G009475

This paper presents flight results for optical angles-only navigation of a satellite swarm, conducted during the Starling Formation-Flying Optical Experiment (StarFOX). StarFOX is a payload of the NASA Starling mission, which consists of four propulsive CubeSats launched in 2023. Angles-only methods apply inter-satellite bearing angles obtained by onboard cameras for navigation. StarFOX flight data provides the first in-flight demonstrations of 1) multi-target and multi-observer angles-only navigation, 2) autonomous onboard initialization of relative navigation for an unknown target, 3) long-term maneuver-free convergence of angles-only orbit estimates, and 4) simultaneous absolute and relative orbit determination using angles-only measurements. StarFOX applies the Absolute and Relative Trajectory Measurement System (ARTMS), which integrates three algorithms. Image processing detects and tracks multiple targets in images, using hypothesis methods and kinematic modeling, and computes target bearing angles. Batch orbit determination computes initial swarm orbit estimates from bearing angle batches via iterative batch least squares and sampling of weakly observable target range. Sequential orbit determination leverages an adaptive, efficient unscented Kalman filter with nonlinear models to refine swarm state estimates over time. Multi-observer measurements shared over an intersatellite link are seamlessly fused to enable robust absolute and relative orbit determination. Relative positioning uncertainties of 1.3% of the target range (1σ) are achieved for a single observer under challenging conditions, reduced to 0.6% (1σ) with multiple observers. Results demonstrate promising performance, robustness, flexibility, and autonomy with regard to enabling future distributed missions.

I. Introduction

DISTRIBUTED Space Systems (DSS) can offer many advantages when compared to traditional monolithic spacecraft, including improved accuracy, coverage, flexibility, robustness, and the ability to achieve entirely new objectives [1]. This has led to the deployment of a variety of DSS science missions [2–5] and the proposed application of DSS to areas such as space situational awareness (SSA) [6] and in-space servicing, assembly, and manufacturing (ISAM) [7]. However, robust navigation for DSS remains a technological challenge. The majority of DSS have been deployed in Earth orbit. Their navigation assumes availability of external metrologies such as Global Navigation Satellite System (GNSS) signals and frequent contact with the ground. Systems outside of Earth orbit may navigate via the Deep Space Network (DSN) or similar resources, but such methods impact timely decision-making for missions and are not easily scalable to future DSS. Furthermore, navigation for non-cooperative objects such as space debris cannot be performed with GNSS. It is therefore necessary to develop new self-contained navigation systems to enable DSS in more varied scenarios, characterized by a high degree of autonomy and robustness. Minimal technical and financial costs for associated hardware are preferred so that miniaturized technology can be leveraged.

Angles-only navigation, in which observer spacecraft obtain bearing angles to target space objects using onboard vision-based sensors (VBS), is a compelling technology in this context. VBS are ubiquitous on modern spacecraft in the form of star trackers. These sensors are passive and inexpensive with a high dynamic range and require minimal mass, volume, and power budgets [8]. If DSS observers are equipped with an inter-satellite link (ISL), as can be implemented with typical radio frequency hardware, measurements from multiple observers can be fused to improve navigation performance. Accordingly, angles-only navigation generally requires no additional hardware even when used on small and inexpensive spacecraft. A further benefit is that optical sensors can obtain measurements of non-cooperative targets. However, bearing angles do not provide explicit target range information, which produces challenging observability conditions [9–12]. Target range is only weakly observable, and it is often difficult to simultaneously estimate absolute and relative orbits [13].

These aspects have motivated extensive research and several prior attempts to conduct angles-only navigation in flight. These include the Orbital Express mission [14], the Advanced Rendezvous using GPS and Optical Navigation (ARGON) experiment [15,16], and the Autonomous Vision Approach Navigation and Target Identification (AVANTI) experiment [17,18]. Of these, ARGON and AVANTI are well-documented in literature.

ARGON (2012) was a key activity of the PRISMA mission, an in-orbit testbed for formation-flying and rendezvous technology developed by several European space centers. A single-observer satellite obtained images of a noncooperative target satellite in low Earth orbit (LEO) during an approach from 30 to 3 km of inter-satellite distance (ISD). Navigation and control tasks were performed in a ground-in-the-loop manner using downlinked images. Target bearing angles were provided to a batch least-squares algorithm, which estimated the target's relative orbit elements over the measurement period. This estimate was passed to a maneuver planner, which generated guidance profiles to both achieve the rendezvous goal and improve range estimation performance. The ground-generated plan was then uplinked.

AVANTI (2016) was conducted during the FIREBIRD mission of the German Aerospace Center (DLR). It encompassed the rendezvous of the BIROS microsatellite and an ejected nanosatellite from 13 km to 50 m of ISD. In comparison to ARGON, AVANTI was more autonomous. Targets could be autonomously tracked in images with less reliance on orbit knowledge from two-line elements, an extended Kalman filter was used to perform relative state estimation on board, and maneuver planning employed a model predictive control strategy for optimization. Passive safety was also enforced via eccentricity/inclination-vector separation. Non-ideal conditions resulted in significant bearing angle measurement

---

Received 12 July 2025; accepted for publication 27 October 2025; published online 4 February 2026.

*Postdoctoral Scholar, Department of Aeronautics and Astronautics, 496 Lomita Mall. jjkruger@stanford.edu. Young Professional AIAA.

†Associate Professor, Department of Aeronautics and Astronautics, 496 Lomita Mall. Associate Fellow AIAA.

outages and degradation of attitude determination performance at closer ranges when target brightness outshone background stars in the field of view (FOV).

The impact of ARGON and AVANTI is felt in their successes and lessons learned. However, they are equivalently characterized by major limitations. Both systems did not consider multiple observers or targets and focused on relative state estimation only, relying on external knowledge of the observer's absolute orbit (e.g., from GNSS) to maintain absolute state convergence. Additional reliance on accurate a priori relative orbit information from the ground to initialize navigation, as well as frequent translational maneuvers to resolve the weakly observable ISD, reduced overall autonomy and efficiency.

The goal of the Starling Formation-Flying Optical Experiment (StarFOX) is to remove these limitations entirely. StarFOX is one of four experimental payloads of the Starling mission managed by the NASA Ames Research Center [19,20]. Starling was launched in July 2023 and is visualized in Fig. 1. It consists of four propulsive 6U CubeSats in LEO and aims to increase the readiness of four enabling technologies for autonomous spacecraft swarms: decision-making, networking, maneuver planning, and absolute and relative navigation. Specifically, StarFOX applies the angles-only Absolute and Relative Trajectory Measurement System (ARTMS). ARTMS is a self-contained architecture [21,22] that provides distributed, autonomous, scalable navigation for DSS orbiting an arbitrary central body, without reliance on maneuvers or external measurement sources. A series of StarFOX experiments has explored ARTMS' flexibility for angles-only navigation in single-observer and distributed multi-observer scenarios, using either ground-assisted or autonomous state initializations, with or without GNSS availability, in two swarm formations. The experiment campaign consequently builds upon ARGON and AVANTI for a more wide-ranging and ambitious demonstration of angles-only navigation in orbit.

The analyses of flight results from StarFOX are the first demonstrations of autonomous angles-only navigation for a satellite swarm [23,24]. Results encompass multi-target and multi-observer relative navigation, autonomous initialization of navigation for unknown targets, and simultaneous absolute and relative orbit determination. Table 1 delineates advances over the state-of-the-art. Results are obtained from 1) ARTMS operating on board the Starling swarm and 2) postprocessing of flight imagery and sensor data using ARTMS flight code running on the ground, with comparisons to preflight expectations. Though significant challenges were encountered in practice, all of StarFOX's experimental objectives were achieved. It thus provides valuable lessons learned for the design and implementation of angles-only navigation on future DSS.

Following this introduction, the angles-only navigation problem is defined in Sec. II. ARTMS and its three constituent algorithms for image processing and orbit determination are described in detail in Sec. III. An overview of the Starling mission and flight operations is presented in Sec. IV. Section V presents quantitative and qualitative analysis of flight results across key scenarios, followed by conclusions in Sec. VI.

Fig. 1 An artist's impression of the Starling swarm. Credit: Blue Canyon Technologies/NASA.

Table 1 Comparison of prior angles-only flight experiments

Parameter	ARGON	AVANTI	StarFOX
Angles-only relative nav.	✓	✓	✓
Angles-only absolute nav.	✗	✗	✓
Estimates computed onboard	✗	✓	✓
Estimates initialized onboard	✗	✗	✓
Maneuver-free convergence	✗	✗	✓
Multi-target navigation	✗	✗	✓
Multi-observer navigation	✗	✗	✓

II. Mathematical Background

A. Measurement Model

ARTMS produces angles-only measurements by computing the time-tagged bearing angles to objects detected in images. Two rotating coordinate frames are defined. First, consider the radial/tangential/normal (RTN) frame of an observer, denoted $\mathcal{R}$. It is centered on and rotates with the observer and consists of orthogonal basis vectors $\hat{\boldsymbol{x}}^\mathcal{R}$ (directed along the observer's absolute position vector); $\hat{\boldsymbol{z}}^\mathcal{R}$ (directed along the observer's orbital angular momentum vector); and $\hat{\boldsymbol{y}}^\mathcal{R} = \hat{\boldsymbol{z}}^\mathcal{R} \times \hat{\boldsymbol{x}}^\mathcal{R}$ [25]. Similarly, define a frame $\mathcal{W}$ using $\hat{\boldsymbol{y}}^\mathcal{W}$ (directed along the observer's velocity vector); $\hat{\boldsymbol{z}}^\mathcal{W} = \hat{\boldsymbol{z}}^\mathcal{R}$; and $\hat{\boldsymbol{x}}^\mathcal{W} = \hat{\boldsymbol{y}}^\mathcal{W} \times \hat{\boldsymbol{z}}^\mathcal{W}$. $\mathcal{W}$ only differs from $\mathcal{R}$ by a rotation of the observer flight path angle ϕ_f about $\hat{\boldsymbol{z}}^\mathcal{R}$ with $\phi_f \approx 0$ in near-circular orbits [25]. Finally, define the observer vision-based sensor (VBS) coordinate frame $\mathcal{V}$ consisting of orthogonal basis vectors $\hat{\boldsymbol{x}}^\mathcal{V}, \hat{\boldsymbol{y}}^\mathcal{V}, \hat{\boldsymbol{z}}^\mathcal{V}$, where $\hat{\boldsymbol{z}}^\mathcal{V} = \hat{\boldsymbol{x}}^\mathcal{V} \times \hat{\boldsymbol{y}}^\mathcal{V}$ is aligned with the camera boresight. The VBS may be pointed as necessary to keep targets in the FOV. In the case of StarFOX, it is chosen to always point the camera boresight in the velocity or antivelocity direction for simplicity, such that $\hat{\boldsymbol{z}}^\mathcal{V}$ is aligned with $\pm\hat{\boldsymbol{y}}^\mathcal{W}$. This was sufficient to image the swarm while avoiding imaging the limb of the Earth. Figure 2 illustrates this scenario.

Bearing angles consist of azimuth and elevation $[\alpha, \epsilon]^\top$. Geometrically, they subtend the line-of-sight vector $\delta\boldsymbol{r}^\mathcal{V}$ from observer to target. They may also be considered as an analog for the image frame coordinates to the target object pixel cluster. The measurement model $\boldsymbol{y}$ for bearing angles from observer o to target t is [21]

$$\delta\boldsymbol{r}^\mathcal{V} = \boldsymbol{r}_t^\mathcal{V} - \boldsymbol{r}_o^\mathcal{V} = [\delta r_x^\mathcal{V}, \delta r_y^\mathcal{V}, \delta r_z^\mathcal{V}]^\top \tag{1}$$

$$\boldsymbol{y}^\mathcal{V}(\delta\boldsymbol{r}^\mathcal{V}) = \begin{bmatrix}\alpha \\ \epsilon\end{bmatrix}^\mathcal{V} = \begin{bmatrix}\arcsin\left(\delta r_y^\mathcal{V}/\|\delta\boldsymbol{r}^\mathcal{V}\|_2\right) \\ \arctan\left(\delta r_x^\mathcal{V}/\delta r_z^\mathcal{V}\right)\end{bmatrix} \tag{2}$$

Measurements and states are also referenced with respect to an inertial reference frame centered on an arbitrary central body, denoted $\mathcal{I}$. The inertial frame for StarFOX is the Earth-centered inertial (ECI) J2000 frame. Bearing angles are related to $\mathcal{I}$ by rotating $\delta\boldsymbol{r}^\mathcal{V}$ into $\mathcal{I}$, as per $\delta\boldsymbol{r}^\mathcal{I} = {}^\mathcal{V}\vec{\boldsymbol{R}}^\mathcal{I}\delta\boldsymbol{r}^\mathcal{V}$, where ${}^\mathcal{V}\vec{\boldsymbol{R}}^\mathcal{I}$ denotes a rotation from frame $\mathcal{V}$ into frame $\mathcal{I}$. This rotation is generally computed by performing attitude determination using stars identified by the VBS [15]. Other relevant rotations ${}^\mathcal{R}\vec{\boldsymbol{R}}^\mathcal{I}$ and ${}^\mathcal{W}\vec{\boldsymbol{R}}^\mathcal{I}$ can be computed using the observer's

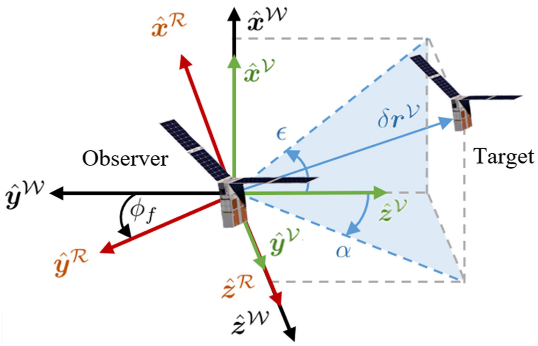

Fig. 2 Definition of coordinate frames and bearing angles with VBS pointing in the $-\hat{y}^\mathcal{W}$ direction.

absolute orbit estimate, which is assumed to be coarsely known. In practice, VBS measurement availability is affected by optical visibility constraints such as eclipse periods or attitude maneuvers by the observer.

B. System State

Because angles-only navigation is characterized by weak observability [9], proper selection of the state parameterization is crucial to maximize estimation accuracy and robustness. ARTMS leverages quasi-nonsingular absolute orbit elements (OEs) and relative orbit elements (ROEs). The absolute orbit $\boldsymbol{\alpha}$ is defined as

$$\begin{aligned}\boldsymbol{\alpha} &= \begin{bmatrix} a & e_x & e_y & i & \Omega & u \end{bmatrix}^\top \\ &= \begin{bmatrix} a & e\cos\omega & e\sin\omega & i & \Omega & \omega + M \end{bmatrix}^\top \end{aligned} \tag{3}$$

where a, e, i, Ω, ω, and M are the Keplerian orbit elements. The relative orbit of each target detected by the onboard sensor, denoted $\delta\boldsymbol{\alpha}$, is described by the quasi-nonsingular ROEs adopted by D'Amico [26]. Each of these ROEs is a function of the orbit elements of t and o, given by

$$\delta\boldsymbol{\alpha} = \begin{pmatrix} \delta a \\ \delta\lambda \\ \delta e_x \\ \delta e_y \\ \delta i_x \\ \delta i_y \end{pmatrix} = \begin{pmatrix} \delta a \\ \delta\lambda \\ |\delta\boldsymbol{e}|\cos\phi \\ |\delta\boldsymbol{e}|\sin\phi \\ |\delta\boldsymbol{i}|\cos\theta \\ |\delta\boldsymbol{i}|\sin\theta \end{pmatrix} = \begin{pmatrix} (a_t - a_o)/a_o \\ (u_t - u_o) + (\Omega_t - \Omega_o)\cos i_o \\ e_{x,t} - e_{x,o} \\ e_{y,t} - e_{y,o} \\ i_t - i_o \\ (\Omega_t - \Omega_o)\sin i_o \end{pmatrix} \tag{4}$$

where $[\delta e_x, \delta e_y]$ are components of the relative eccentricity vector with phase ϕ and $[\delta i_x, \delta i_y]$ are components of the relative inclination vector with phase θ.

These state definitions have been previously studied in literature, resulting in the development of accurate analytical dynamics models [27,28]. Also, these states are slowly varying and enable accurate numerical integration using Gauss's variational equations (GVE) with large time steps for efficient onboard orbit propagation [29]. More importantly, for loose satellite formations and swarms in which the inter-satellite separation is primarily along-track (e.g., string-of-pearls formations), the weakly observable range to each target is primarily captured by the $\delta\lambda$ term. This allows ARTMS to maximize accuracy by applying separate state estimation techniques to different state components. Additionally, it has been shown that the semimajor axis of the observer's orbit is strongly observable using bearing angle measurements to a single target [13]. Combined, these properties enable accurate and computationally efficient estimation algorithms with minimal reliance on a priori information. The complete estimated state vector $\boldsymbol{x}$ on board each spacecraft was

$$\boldsymbol{x} = \begin{bmatrix} \boldsymbol{\alpha} & \boldsymbol{d}^{\mathcal{R}} & \delta\boldsymbol{\alpha}_1 & \delta\boldsymbol{d}_1^{\mathcal{R}} & \delta\boldsymbol{\alpha}_2 & \delta\boldsymbol{d}_2^{\mathcal{R}} & \delta\boldsymbol{\alpha}_3 & \delta\boldsymbol{d}_3^{\mathcal{R}} \end{bmatrix} \in \mathbb{R}^{36} \tag{5}$$

which encompasses observer OEs, the ROEs of up to three swarm targets, as well as $\boldsymbol{d}^{\mathcal{R}}$ and $\delta\boldsymbol{d}^{\mathcal{R}}$. These terms are the empirical estimates of absolute and relative perturbing accelerations acting on the swarm in $\mathcal{R}$, defined as

$$\boldsymbol{d}^{\mathcal{R}} = \begin{bmatrix} d_x^{\mathcal{R}} & d_y^{\mathcal{R}} & d_z^{\mathcal{R}} \end{bmatrix}^\top \tag{6}$$

The state can be extended to include an arbitrary number of targets but was limited to three in flight, as this was the maximum number of targets within the Starling swarm. ARTMS also possesses the capacity to estimate auxiliary state components, such as differential clock offsets and drift rates between observers, differential ballistic coefficients between the observer and targets, and observer sensor biases [11,21,30,31]. Only orbit estimation was performed during the initial StarFOX experiment phase, to reduce computation costs and system complexity.

C. Dynamics Model

ARTMS propagates absolute orbits using fourth-order Runge–Kutta integration of the GVE. For state $\boldsymbol{\alpha}$, the osculating OEs of each spacecraft evolve according to

$$\dot{\boldsymbol{\alpha}} = G(\boldsymbol{\alpha})\boldsymbol{d}^{\mathcal{R}} \tag{7}$$

where $G \in \mathbb{R}^{6\times3}$ is the GVE matrix [32]. Depending on the orbit regime, common perturbations include spherical harmonic gravity terms, atmospheric drag, third-body gravity, and solar radiation pressure [33]. In flight, ARTMS typically applies a 10×10 GGM01S [2] spherical harmonic gravity model, a Harris–Priester atmosphere model with cannonball drag ($B_D \approx 0.015$), and a 30 s RK4 step. Analytic dynamics models for the mean OEs, which include the effects of J_2 gravity [34] are used within ARTMS when computational efficiency is paramount, e.g., during multi-hypothesis tracking and sample-based batch orbit determination, as outlined in Sec. III.

With regard to relative dynamics, the ROEs provide useful geometric intuition regarding relative motion. The curvilinear position vector of a target in the observer's RTN frame is defined as $\delta\boldsymbol{r}_{\text{curv}}^{\mathcal{R}} = (\delta r, a_o\beta, a_o\gamma)$ [35]. Here, $\delta r, \beta, \gamma$ are differences in orbit radii, angular in-plane separations, and angular out-of-plane separations, respectively. The curvilinear representation captures the effects of orbit curvature with improved accuracy and can be mapped back to rectilinear coordinates via

$$\delta\boldsymbol{r}_{\text{rect}}^{\mathcal{R}} = \begin{bmatrix} (a_o + \delta r)\cos\beta\cos\gamma - a_o \\ (a_o + \delta r)\sin\beta\cos\gamma \\ (a_o + \delta r)\sin\gamma \end{bmatrix} \tag{8}$$

As first shown by D'Amico [26], there exists a linear mapping between the ROEs and $\delta\boldsymbol{r}_{\text{curv}}^{\mathcal{R}}$, defined as

$$\delta\boldsymbol{r}^{\mathcal{R}} \approx a_o \begin{bmatrix} \delta a - \delta e_x\cos u_o \\ -1.5\delta a + \delta\lambda + 2\delta e_x \sin u_o - 2\delta e_y \cos u_o \\ \delta i_x \sin u_o - \delta i_y \cos u_o \end{bmatrix} \tag{9}$$

The mapping was later extended to eccentric orbits by defining the eccentric ROEs [36] $\delta\boldsymbol{\alpha}^* = [\delta a, \delta\lambda^*, \delta e_x^*, \delta e_y^*, \delta i_x, \delta i_y]^\top$, which revert to the quasi-nonsingular ROEs for $e_o \approx 0$.

Figure 3 presents relative motion in RTN for this mapping for small ISD. Components of oscillatory motion produced by the

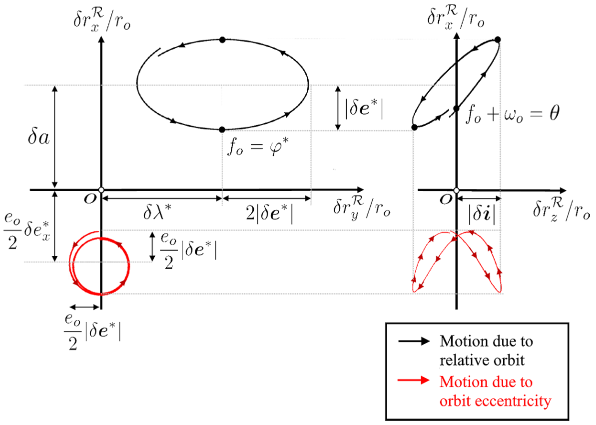

Fig. 3 Components of relative motion in the $\hat{x}^{\mathcal{R}} - \hat{y}^{\mathcal{R}}$ (RT) and $\hat{x}^{\mathcal{R}} - \hat{z}^{\mathcal{R}}$ (RN) planes.

target's separation and relative orbit are shown in black, possessing the same frequency as the orbit. Components of oscillatory motion produced by orbit eccentricity are shown in red, acting at twice the frequency of the orbit. Note that δa and $\delta\lambda^*$ capture mean offsets in the radial and along-track directions, respectively; magnitudes of δe^* and δi correspond to magnitudes of oscillations in the RT and RN planes, respectively; and phases of δe^* and δi dictate the orientation and aspect ratio of the tilted ellipse in the RN plane. The eccentricity of the observer's orbit superimposes additional offsets and higher-frequency oscillations in the RT and RN planes.

III. Navigation Algorithms

To simplify discussions, the following terminologies are adopted. The "observer" refers to the spacecraft hosting the instance of ARTMS in question. A "remote observer" is another spacecraft hosting ARTMS, providing measurements over the inter-satellite link (ISL). The "local subsystem" includes an observer and all its "targets," which are the resident space objects (RSO) detected by the onboard VBS. The "system" refers to the entire distributed space system (DSS), consisting of all involved observers and targets. Targets may be non-cooperative objects that do not actively assist navigation or may be cooperative remote observers themselves. Figure 4 presents a notional illustration of a four-spacecraft DSS using ARTMS.

A high-level overview of ARTMS as implemented in the StarFOX flight software is shown in Fig. 5. The architecture consists of three core modules based on algorithms recently developed at Stanford's Space Rendezvous Laboratory: Image Processing (IMP) [37], Batch Orbit Determination (BOD) [13], and Sequential Orbit Determination (SOD) [21]. The VBS provides time-tagged raw images, which are processed to obtain inertial bearing angles to RSO. The ISL communicates orbit estimates and angle measurements between observers in the DSS, which allows ARTMS to perform distributed multi-observer navigation. The ground segment provides telecommands, maneuver plans, and orbit estimates to each observer in the DSS, and it receives ARTMS telemetry. This constitutes the only a priori information needed, i.e., each observer must possess a coarse estimate of its own orbit at a single epoch to initialize. If the spacecraft is equipped with a GNSS receiver, it optionally provides position/velocity/time (PVT) navigation solutions. This may assist the observer in maintaining a more timely and accurate absolute orbit estimate when necessary.

The IMP module uses a coarse estimate of the observer's orbit and VBS images to produce batches of bearing angles and corresponding uncertainties to all visible targets without a priori relative orbit knowledge. Relative orbit information from SOD can be leveraged when available to reduce computation cost. IMP measurement batches are provided to the BOD and SOD modules. Additionally, IMP sends the orbit estimate and bearing angles to the ISL for transmission.

The BOD module uses the orbit estimate and angle batch from IMP to compute orbit estimates for all spacecraft in the local system (including itself and all targets observed by onboard cameras). This DSS state estimate is provided to SOD for initialization and fault detection.

The SOD module uses the state estimate from BOD to initialize a navigation filter that continuously estimates the orbits of all spacecraft in the local system as well as auxiliary parameters (e.g., empirical accelerations). SOD seamlessly fuses measurements from IMP and from remote observers communicated over the ISL. SOD state estimates are provided to the ground and also to IMP to assist with tracking targets.

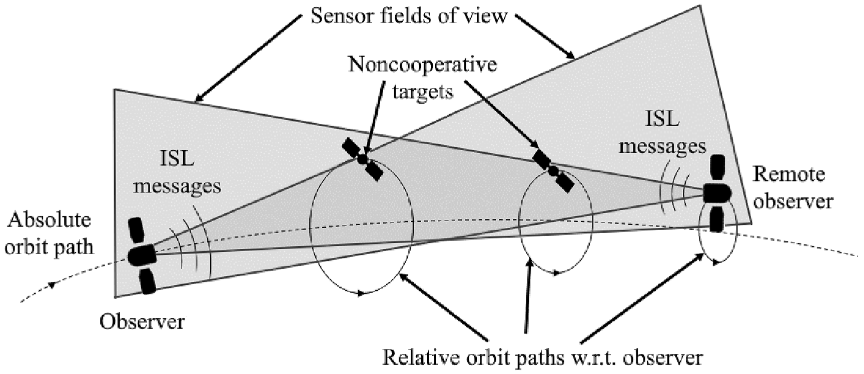

Fig. 4 Notional illustration of ARTMS observers and targets for a four-spacecraft system (not to scale).

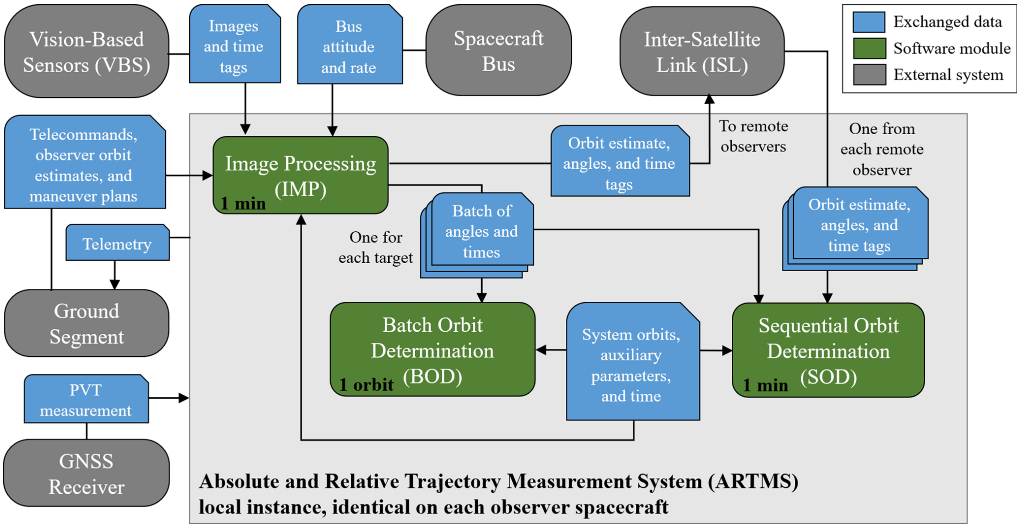

Fig. 5 A high-level overview of the ARTMS flight architecture.

Overall, ARTMS provides real-time orbit estimates for the host spacecraft and each target detected by the onboard sensor. The only hardware requirements posed on the spacecraft are that it must have a VBS and ISL. Navigation is primarily autonomous in that minimal a priori information is required and no external measurement sources are needed. The ISL enables cooperative navigation, and the system is distributed in the sense that each observer only navigates for its local subsystem, which may be a subset of the complete DSS.

A. Image Processing

The objective of the IMP module is to produce batches of time-tagged bearing angle measurements for each target, using images from the onboard VBS. This is accomplished in two phases. First, IMP processes incoming images to reduce them to a set of inertial bearing angles that may correspond to resident space objects. Second, known targets are tracked and new targets are detected using an approach inspired by multi-hypothesis tracking (MHT) [38]. Figure 6 presents an example image input. For StarFOX operating in LEO, typical sample rates were 60–120 s between images, equating to 50–100 images per orbit. Image integration times were approximately 75 ms.

The first phase of IMP uses a set of state-of-the-art algorithms. An iterative weighted center of gravity [39] or fast Gaussian fitting algorithm [40] is used to simplify the raw image into a list of pixel cluster centroids. Centroids are converted to line-of-sight (LOS) unit vectors in the VBS frame using a calibrated sensor model. Next, the Pyramid star identification algorithm [41] is applied to remove stars from the list of vectors. Uncatalogued stars are detected by considering objects with unchanging inertial unit vectors between images. Similarly, camera hotspots are removed by considering objects with unchanging pixel coordinates. The VBS attitude is estimated via the q-method, applied to the LOS vectors to identify stars in the inertial and VBS reference frames [42]. Attitude is computed using the same sensor that provides the bearing angle measurements to minimize calibration errors. The remaining minimalistic set of inertial unit vectors likely corresponds to known targets or other unknown objects in the FOV.

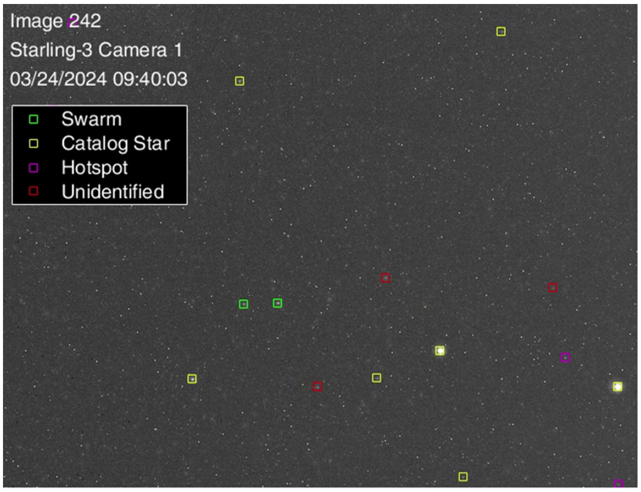

Fig. 6 Part of a Starling flight image with point sources detected and identified by IMP.

In the second phase, these measurement candidates must be consistently assigned to new or tracked targets without a priori relative orbit knowledge. To accomplish this, IMP employs the Spacecraft Angles-Only Multi-Target Tracking System (SAMUS) [37]. SAMUS only requires a coarse estimate of the observer's absolute orbit and that targets do not perform unknown translational maneuvers during tracking. SAMUS is valid for eccentric orbits and meets the constraints of risk-averse angles-only navigation in space; i.e., it achieves close to 100% measurement assignment precision with low measurement frequencies and limited computational resources.

SAMUS applies the core concept of MHT in that as measurements arrive, many simultaneous hypotheses are maintained as to how they can be associated into target tracks. The algorithm converges toward the correct hypothesis over time as more measurement information becomes available. MHT is chosen as a basis because it is mature and demonstrably accurate, with its primary disadvantage being the need to frequently and heuristically prune hypotheses for real-time computation [43]. To overcome this, SAMUS applies domain-specific knowledge to develop precise pruning criteria. Figure 7 presents an overview of core SAMUS operations, where dashed outlines denote steps that only occur at relevant epochs.

The first operation initializes possible new tracks by applying the density-based spatial clustering of applications with noise (DBSCAN) algorithm [44] to the merged set of unidentified measurements from the most recent four images. During StarFOX and similar swarm scenarios, target velocities are low compared to other objects in the FOV because they are in similar absolute orbits to observers. Thus, their measurements form clusters in $\mathcal{V}$. A new SAMUS target is initialized if DBSCAN detects a cluster of $n_D \geq 4$ unidentified bearing angles within some small radius ϵ_D.

Subsequent operations apply Eq. (9), which maps the OEs and ROEs to a target's curvilinear position vector in the observer's RTN frame. On the right-hand side of Eq. (9), the mean argument of latitude u_o is the only quickly varying term, whereas all other terms vary slowly in the presence of perturbations [26]. The second step in

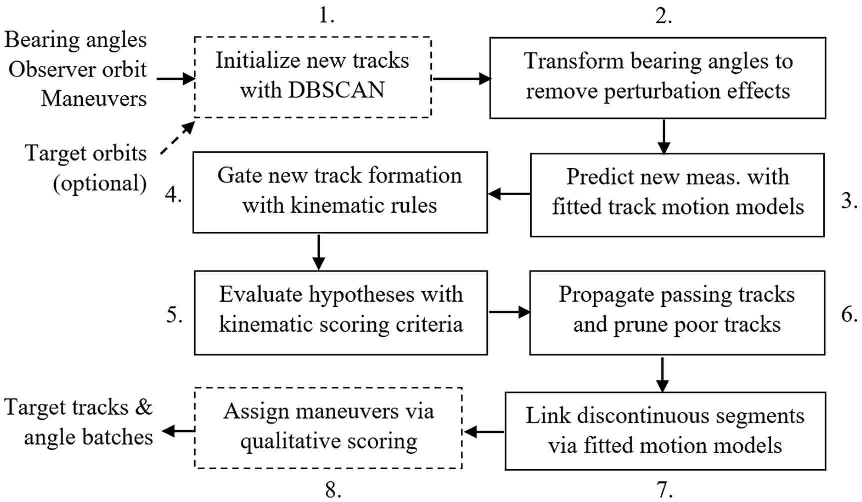

Fig. 7 SAMUS algorithm summary and core sequence of operations.

Fig. 7 recalls that system members in swarm scenarios remain in close proximity in inertial space. Each is therefore affected similarly by perturbations. By synchronously differencing the bearing angles of different targets (using one track as a virtual, moving origin for another), perturbation effects are approximately canceled between targets, recovering motion as per Eq. (9). Relative motion is then described by periodic functions with known form, parametric in u_o. Even if the specific ROEs are unknown, this form provides expectations that can be leveraged to assess bearing angle tracks.

The third operation rearranges the radial and normal components of Eq. (9) into separable linear systems in azimuth and elevation. Given past measurements in a track and corresponding estimates of the observer's orbit, a motion model may be fitted to the track, as per

$$\begin{bmatrix} \epsilon \\ \alpha \end{bmatrix}^{\mathcal{R}} \approx \frac{r_o}{a_o} \begin{bmatrix} x_1 - x_2 \cos(u_o - x_3) \\ x_4 + x_5 \sin(u_o - x_6) \end{bmatrix} \tag{10}$$

where $x_{1,\ldots,6}$ are scaled ROE equivalents in bearing angle space [37]. Terms r_o, a_o, and u_o are computed from the absolute orbit estimate. If a track consists of at least three bearing angles at different times, the six unknowns $x_{1,\ldots,6}$ can be solved for via least squares. Measurements in future epochs can then be predicted by propagating the absolute orbit estimate and applying the fitted model. Goodness of fit can be assessed via residuals.

In the fourth operation, SAMUS applies a set of kinematic rules [37] derived from the parametric motion model to assess which tracks are physically reasonable. Figure 8 presents relevant track quantities in the two-dimensional bearing angle plane. Vectors $\Delta\vec{y}_k$ connect successive bearing angles y_k, where the subscript k refers to the current tracking epoch. Each $\Delta\vec{y}_k$ has magnitude d_k, phase ζ_k, and inter-vector angle ψ_k. Briefly, the rules are summarized as follows:

1) Track velocities in the bearing angle plane must be below a maximum, as derived from orbit dynamics, the image measurement frequency, and constraints on the allowed magnitude of relative motion such that $d_k < d_{k,\max}$.

2) Track velocities must be consistent over time such that their ratio between successive images is close to one such that $|d_k/d_{k-1}| \approx 1$.

3) Tracks should turn in a gradual manner as per the relative motion ellipse such that $\psi_k < \psi_{k,\max}$.

4) Tracks should turn in a consistent direction as per the relative motion ellipse such that $\text{sign}(\zeta_k - \zeta_{k-1}) = \text{sign}(\zeta_{k-1} - \zeta_{k-2})$.

5) New track measurements must be close to the measurement predicted by the fitted motion model within some tolerance, as per $\|\boldsymbol{y}_k - \boldsymbol{y}_k^p\|_2 < d_{p,\max}$.

Only tracks that pass all rules are propagated. Their application greatly increases the efficiency of MHT by preventing the formation of unlikely tracks.

When multiple tracks pass all rules, SAMUS scores propagated tracks via a set of criteria that assess how well each fulfills the expectations of Eq. (9) and their prior motion. These criteria include the following:

1) The size of summed residuals from track fitting

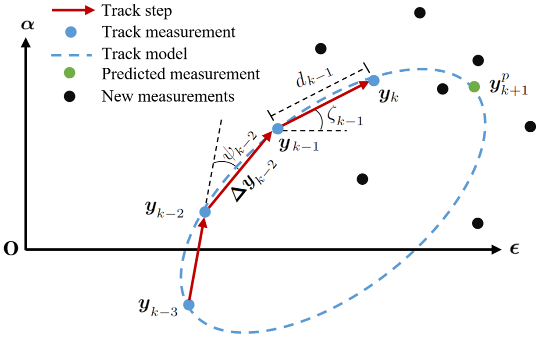

Fig. 8 Target track quantities in the two-dimensional bearing angle plane.

2) The difference between predicted and assigned measurements, i.e., $\|\boldsymbol{y}_k - \boldsymbol{y}_p\|_2$

3) The closeness of d_k to its predicted value and mean values, i.e., $|d_k - d_p|$ and $|d_k - d_{\text{mean}}|$

4) The closeness of ψ_k to its predicted value and mean values, i.e., $|\psi_k - \psi_p|$ and $|\psi_k - \psi_{\text{mean}}|$

5) A preference for tracks to possess lower velocities and more gradual turns, i.e., $|d_k|$ and $|\psi_k|$

Lower scores are preferred. In contrast to traditional MHT methods, which often rely on a single Mahalanobis distance metric for scoring, SAMUS aims to be more robust. Often, tracks intersect or are in close proximity in the image plane, or motion between images is on the order of VBS noise. A single scoring metric is therefore not robust. By using a larger set of metrics, consensus supports the correct choice over time, even if some metrics temporarily favor incorrect hypotheses. Additionally, scoring does not require probabilistic estimates of false alarm densities or target decay rates, which are not easily obtainable for spacecraft.

In the sixth operation, surviving hypotheses are managed and pruned according to traditional MHT methods [43]. If a track is not observed in an image (e.g., because measurement availability is degraded), it may be propagated using the predicted track measurement as a fallback. The maximum number of missed observations before pruning is defined by the user.

The seventh operation provides a re-initialization method for tracks in the event of longer interruptions such as eclipse periods. To connect shorter tracks on either side of an interruption, the aforementioned linear system fit is computed for every possible set of track pairs. The combination of compatible pairs that produces the least fitting residuals is chosen as output. Finally, if maneuvers are expected during tracking, they are matched to tracks in the eighth operation by examining the qualitative change in their fitted $x_{1,\ldots,6}$ values pre- and postmaneuver. This is compared to the expected qualitative changes in ROEs from a state transition matrix.

SAMUS is also able to apply target state knowledge from SOD (if available) to reduce computation cost. Relative orbit estimates are propagated into the image epoch to provide predicted measurements. The kinematic rules are replaced by a validity region around the prediction, computed via an unscented transform of the state covariance into bearing angle space. The Mahalanobis distance [45] between predicted and assigned measurement is employed for scoring, defined as

$$\sqrt{(\boldsymbol{y}_k^p - \boldsymbol{y}_k)^\top (\boldsymbol{\Sigma}_k^p)^{-1} (\boldsymbol{y}_k^p - \boldsymbol{y}_k)} \tag{11}$$

where $\boldsymbol{\Sigma}_k^p$ is the bearing angle covariance.

B. Batch Orbit Determination

The BOD module must produce orbit estimates for the local swarm with sufficient accuracy to initialize the SOD module, using only a coarse estimate of the observer's orbit and batches of bearing angles from IMP. Nominally, the module operates on measurement batches of [50, 200] measurements across 1–2 orbits per target. State estimation is accomplished using an algorithm [13] inspired by [46] that separately estimates the relative orbits of each target while refining the observer's semimajor axis estimate. The batch estimation problem is somewhat ill-conditioned due to the weak angles-only observability of $\delta\lambda$. In typical swarm or formation-flying scenarios with large in-track separations, such as StarFOX, $\delta\lambda$ most closely corresponds to the inter-satellite range. As a result, $\delta\lambda$ is estimated using a line sampling procedure, whereas the other strongly observable ROEs are estimated using a Gauss–Newton algorithm (iterative batch least squares).

Consider a model $\boldsymbol{y}^{\mathcal{I}}(\boldsymbol{x}(t_{\text{est}}), t)$ that provides bearing angles at time t as a function of system state $\boldsymbol{x}(t_{\text{est}})$ at estimation epoch t_{est}. For each target, the system state is partitioned into estimated components and components provided a priori, with $\boldsymbol{x} = [\boldsymbol{x}_{\text{est}}^\top, \boldsymbol{x}_{\text{prior}}^\top]^\top$ and $\boldsymbol{x}_{\text{est}} \in \mathbb{R}^{N_x}$. StarFOX applies

$$\boldsymbol{x}_{\text{est}} = [a, \delta a, \delta e_x, \delta e_y, \delta i_x, \delta i_y] \tag{12}$$

$$\boldsymbol{x}_{\text{prior}} = [e_x, e_y, i, \Omega, u, \delta\lambda] \tag{13}$$

because the $\boldsymbol{x}_{\text{est}}$ components are more strongly observable [13] than those in $\boldsymbol{x}_{\text{prior}}$. The absolute OEs in $\boldsymbol{x}_{\text{prior}}$ are derived from orbit knowledge at t_{est}, whereas the weakly observable $\delta\lambda$ is specified as a 1-D family of values. The $\delta\lambda$ state space may be bounded by the expected minimum swarm separation and the maximum range physically observable by the VBS. The $\delta\lambda$ sampling interval is selected with reference to onboard computational resources or the desired output accuracy. StarFOX divides the $\delta\lambda$ state space into 100–300 intervals in the positive and negative directions with sampling intervals of 1–2 km.

The BOD module operates on bearing angles provided across N_y epochs $t_1, \ldots, t_{N_y}$, collectively denoted $\boldsymbol{t}$. A measurement batch $\boldsymbol{z}(\boldsymbol{x}, t_{\text{est}}, \boldsymbol{t}) \in \mathbb{R}^{2N_y}$ is

$$\boldsymbol{z}(\boldsymbol{x}(t_{\text{est}}), \boldsymbol{t}) = \begin{pmatrix} \boldsymbol{y}^{\mathcal{I}}(\boldsymbol{x}(t_{\text{est}}), t_1) \\ \vdots \\ \boldsymbol{y}^{\mathcal{I}}(\boldsymbol{x}(t_{\text{est}}), t_{N_y}) \end{pmatrix} \tag{14}$$

Sensitivity matrices containing the partial derivatives of measurements with respect to state components are

$$\boldsymbol{S}_{\text{est}}(\boldsymbol{x}(t_{\text{est}}), \boldsymbol{t}) = \left.\frac{\partial \boldsymbol{z}(\boldsymbol{x}(t_{\text{est}}), \boldsymbol{t})}{\partial \boldsymbol{x}_{\text{est}}}\right|_{\boldsymbol{x}(t_{\text{est}})} \tag{15}$$

$$\boldsymbol{S}_{\text{prior}}(\boldsymbol{x}(t_{\text{est}}), \boldsymbol{t}) = \left.\frac{\partial \boldsymbol{z}(\boldsymbol{x}(t_{\text{est}}), \boldsymbol{t})}{\partial \boldsymbol{x}_{\text{prior}}}\right|_{\boldsymbol{x}(t_{\text{est}})} \tag{16}$$

with $\boldsymbol{S}_{\text{est}} \in \mathbb{R}^{2N_y \times N_x}$. These matrices may be computed via linearization of the measurement model.

BOD initializes the $\boldsymbol{x}_{\text{est}}$ vector for the first $\delta\lambda$ sample by setting all ROEs (except $\delta\lambda$) to zero, obtaining a from absolute orbit knowledge. For subsequent $\delta\lambda$ samples, the ROEs are initialized as

$$\delta\boldsymbol{\alpha}^{+} = \delta\boldsymbol{\alpha}^{-}\frac{\delta\lambda^{+}}{\delta\lambda^{-}} \qquad a^{+} = a^{-} \tag{17}$$

where superscripts "+" and "−" refer to the prior and current sample, respectively. This saves computation because the direction of converged ROE solutions does not change significantly between adjacent $\delta\lambda$ samples.

BOD then performs iterative batch least-squares refinement for each sample. Define the batch of angle measurements from IMP as $\boldsymbol{z}_{\text{meas}}$ and a batch of modeled angle measurements $\boldsymbol{z}_{\text{model}}$ computed by propagating $\boldsymbol{x}_{\text{est}}$ to each epoch in $\boldsymbol{t}$. The difference between measured and modeled angles is

$$\Delta \boldsymbol{z} = \boldsymbol{z}_{\text{meas}} - \boldsymbol{z}_{\text{model}}(\boldsymbol{x}(t_{\text{est}}), \boldsymbol{t}) \tag{18}$$

Subsequently, the state estimate update

$$\Delta \boldsymbol{z} = \boldsymbol{S}_{\text{est}}(\boldsymbol{x}(t_{\text{est}}), \boldsymbol{t}) \Delta \boldsymbol{x}_{\text{est}} \tag{19}$$

$$\boldsymbol{x}_{\text{est}} \leftarrow \boldsymbol{x}_{\text{est}} + \Delta \boldsymbol{x}_{\text{est}} \tag{20}$$

is computed and the unknown $\Delta \boldsymbol{x}_{\text{est}}$ is solved for via least squares. Refinement continues until an iteration limit is reached or $\Delta \boldsymbol{x}_{\text{est}}$ is smaller than a convergence threshold.

BOD assembles the $\boldsymbol{x}_{\text{final}}$ estimate for output from the $\delta\lambda^{*}$ candidate and corresponding converged $\boldsymbol{x}_{\text{est}^{*}}$, which produced the minimum measurement residual. This leads to an optimization objective of

$$\min_{\boldsymbol{x}_{\text{est}}, \delta\lambda} \quad \Delta \boldsymbol{z}^{\top} \Delta \boldsymbol{z} \tag{21}$$

for which $\boldsymbol{x}_{\text{final}} = [\boldsymbol{x}_{\text{prior}}, \boldsymbol{x}_{\text{est}}^{*}, \delta\lambda^{*}]$ is the optimal solution. Figure 9 shows the norms of the converged measurement residual vectors for each candidate value of $\delta\lambda$ in a single test case of the BOD algorithm. The norm of the measurement residual vector is normally a convex function of range, despite its weak observability [46].

It is then necessary to estimate the associated initial state uncertainty. The equation applied is [13]

$$\boldsymbol{R}_{\text{est}} = \boldsymbol{S}_{\text{est}} \boldsymbol{P}_{\text{est}} \boldsymbol{S}_{\text{est}}^{\top} \tag{22}$$

which maps measurement covariance $\boldsymbol{R}_{\text{est}} \in \mathbb{R}^{2N_y \times 2N_y}$ to state covariance $\boldsymbol{P}_{\text{est}} \in \mathbb{R}^{N_x \times N_x}$ via sensitivity $\boldsymbol{S}_{\text{est}}$. To account for all relevant uncertainties, the measurement covariance can be granularly expressed as

$$\begin{aligned} \boldsymbol{R}_{\text{est}} &= \boldsymbol{R}_{\text{prior}} + \boldsymbol{R}_{\text{sens}} + \boldsymbol{R}_{\text{dyn}} = \boldsymbol{S}_{\text{prior}}(\boldsymbol{x}_{\text{final}}) \boldsymbol{P}_{\text{prior}} \boldsymbol{S}_{\text{prior}}(\boldsymbol{x}_{\text{final}}) \\ &\quad + \begin{bmatrix} \boldsymbol{R}_{\text{meas}} & \ldots & \boldsymbol{0} \\ \vdots & \ddots & \vdots \\ \boldsymbol{0} & \ldots & \boldsymbol{R}_{\text{meas}} \end{bmatrix} + \begin{bmatrix} \boldsymbol{R}_{\text{proc}} & \ldots & \boldsymbol{0} \\ \vdots & \ddots & \vdots \\ \boldsymbol{0} & \ldots & \boldsymbol{R}_{\text{proc}} \end{bmatrix} \end{aligned} \tag{23}$$

The $\boldsymbol{R}_{\text{prior}}$ term encompasses uncertainties stemming from a priori information; the $\boldsymbol{R}_{\text{sens}}$ term encompasses uncertainties due to sensor noise; and the $\boldsymbol{R}_{\text{dyn}}$ term encompasses uncertainties due to systematic discrepancies between the onboard and ground truth dynamics models.

The uncertainty of a priori state information, $\boldsymbol{P}_{\text{prior}}$, is assumed to be provided. The sensor noise associated with individual measurements, $\boldsymbol{R}_{\text{meas}} \in \mathbb{R}^{2\times 2}$, is estimated via

$$\begin{aligned} \boldsymbol{R}_{\text{meas}} &= \frac{1}{N_y} \sum_{j=1}^{N_y} \Big((\boldsymbol{z}_{\text{meas}}(t_j) - \boldsymbol{y}^{\mathcal{I}}(\boldsymbol{x}_{\text{final}}, t_{\text{est}}, t_j))(\boldsymbol{z}_{\text{meas}}(t_j) \\ &\quad - \boldsymbol{y}^{\mathcal{I}}(\boldsymbol{x}_{\text{final}}, t_{\text{est}}, t_j))^{\top} \Big) \end{aligned} \tag{24}$$

This estimation of sensor noise using postfit measurement residuals allows operation even when an accurate a priori model of sensor noise is unavailable. The measurement uncertainty due to process noise, $\boldsymbol{R}_{\text{proc}} \in \mathbb{R}^{2\times 2}$, can be estimated by propagating the time-varying orbit covariance due to unmodeled RTN accelerations to each measurement epoch, then performing an unscented transform into bearing angle space [11]. In the cited formulation, the measurement at t_{est} possesses zero dynamics uncertainty (as it requires no

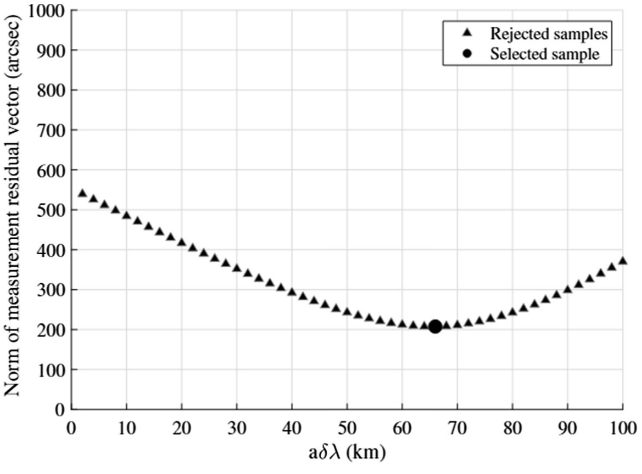

Fig. 9 Behavior of converged measurement residuals for different $\delta\lambda$ samples in the BOD module.

propagation to model), whereas the measurement farthest from t_{est} in time is affected by the most dynamics uncertainty. In this fashion, time-varying dynamics uncertainties are accounted for.

BOD is then able to estimate the initial state uncertainty from known terms via

$$\boldsymbol{P}_{\text{est}} = \boldsymbol{S}^{*}_{\text{est}} \boldsymbol{R}_{\text{est}} \boldsymbol{S}^{*\top}_{\text{est}} \tag{25}$$

where * denotes a pseudoinverse. This formulation allows the BOD module to seamlessly transition between domains where uncertainty is driven by different sources. The complete orbit estimate for the subsystem is formed by concatenating the ROEs in $\boldsymbol{x}_{\text{final}}$ for each target, the mean refined observer semimajor axis, and $\boldsymbol{x}_{\text{prior}}$.

The sampling approach of BOD is enabled by a fully analytical dynamics model, including the Earth oblateness J_2 perturbation [27], which minimizes the computational load of orbit propagation in LEO. Similar analytic models can be employed in other orbit regimes [30,31]. The computation cost of the algorithm increases linearly with the size of the subsystem and the number of $\delta\lambda$ samples, allowing efficient scaling. However, the BOD method developed for StarFOX is less applicable to cases where relative motion is centered on the observer (i.e., $\delta\lambda \approx 0$), because the scale of relative motion is no longer captured by $\delta\lambda$. For these relative motion geometries, other authors have explored alternative methods for angles-only initial relative orbit determination [47,48].

C. Sequential Orbit Determination

The SOD module continually refines estimates of the orbits of all spacecraft in the local swarm as well as auxiliary parameters (e.g., empirical accelerations) by seamlessly fusing measurements from all observers transmitted over the ISL. SOD is based on an adaptive, efficient unscented Kalman filter (UKF) [21], which applies the bearing angle measurement model and numerical GVE dynamics model from Sec. II. In comparison to an extended Kalman filter, the UKF is able to incorporate fully nonlinear dynamics and measurement models and can more accurately approximate higher-order behavior. The resulting observability improvements are critical for enabling maneuver-free state convergence using bearing angles from a single camera, without excessive computational cost [21].

Three additional features are included to maximize performance. First, adaptive process noise estimation is used to improve convergence speed and robustness to errors in the dynamics model [10]. Second, the state definition is organized in a way that exploits the structure of the Cholesky factorization. This reduces the number of calls to the orbit propagator by almost a factor of two [49]. Third, measurement assignment algorithms are used to determine whether measurements received from remote observers are observations of locally tracked targets. This enables distributed stereovision and provides more geometric information to the UKF, with improvements to observability and navigation performance.

The measurement assignment proceeds as follows: SOD first attempts to determine whether any remote observers are in fact targets being tracked by the local ARTMS instance. Let remote observer m have the inertial Cartesian state $\boldsymbol{x}^{\mathcal{I}}_m$ with associated covariance $\boldsymbol{P}^{\mathcal{I}}_m$ (as per its broadcast absolute orbit estimate). Let local target n have the inertial Cartesian state $\boldsymbol{x}^{\mathcal{I}}_n$ with associated covariance $\boldsymbol{P}^{\mathcal{I}}_n$ (as per the swarm state estimate). The Mahalanobis distance σ_{mn} between state estimates is computed as

$$\sigma_{mn} = \sqrt{(\boldsymbol{x}^{\mathcal{I}}_m - \boldsymbol{x}^{\mathcal{I}}_n)^\top (\boldsymbol{P}^{\mathcal{I}}_m + \boldsymbol{P}^{\mathcal{I}}_n)^{-1} (\boldsymbol{x}^{\mathcal{I}}_m - \boldsymbol{x}^{\mathcal{I}}_n)} \tag{26}$$

Remote observer m is identified as local target n if four conditions are fulfilled, as per user-defined thresholds:

1) m has not yet been identified.

2) $\sigma_{mn} \le \sigma_{\text{assign}}$; i.e., the remote observer state is similar to the local target state.

3) $\sigma_{pn} \ge \sigma_{\text{safe}} \ \ \forall \ p \ne m$; i.e., there is no other remote observer state similar to the local target state.

4) $\sigma_{mq} \ge \sigma_{\text{safe}} \ \ \forall \ q \ne n$; i.e., there is no other local target state similar to the remote observer state.

Identifications are removed if $\sigma_{mn} \ge \sigma_{\text{remove}}$, with $\sigma_{\text{remove}} > \sigma_{\text{safe}} > \sigma_{\text{assign}} > 0$. These conditions ensure that targets are formally identified and measurements broadcast by m will not be considered measurements of m itself when they are received by the onboard ARTMS instance (to within a statistical certainty determined by σ values). Preventing this contradiction improves robustness and reduces the search space when assigning measurements.

SOD then attempts to determine whether any measurements from remote observers correspond to targets being tracked by the local ARTMS instance. Let $\boldsymbol{y}^{\mathcal{V}_m}_i$ be bearing angle i received from a remote observer m in its VBS frame. Let $\boldsymbol{y}^{\mathcal{V}_m}_j$ be the modeled bearing angle of object j in the local subsystem (either a target or the local observer) as would be seen by m in its VBS frame. The corresponding bearing angle covariance $\boldsymbol{P}^{\mathcal{V}_m}_j$ is also computed via an unscented transform of the associated state covariances. The Mahalanobis distance σ_{ij} between the measured and modeled angles is then

$$\sigma_{ij} = \sqrt{(\boldsymbol{y}^{\mathcal{V}_m}_i - \boldsymbol{y}^{\mathcal{V}_m}_j)^\top (\boldsymbol{P}^{\mathcal{V}_m}_j)^{-1} (\boldsymbol{y}^{\mathcal{V}_m}_i - \boldsymbol{y}^{\mathcal{V}_m}_j)} \tag{27}$$

To minimize erroneous assignments, measurement i from the remote observer is assigned to local object j if three conditions are satisfied, as per user-defined thresholds:

1) $\sigma_{ij} \le \sigma_{\text{match}}$; i.e., the remote measurement is close to the modeled measurement.

2) $\sigma_{kj} \ge \sigma_{\text{ambig}} \ \ \forall \ k \ne i$; i.e., no other remote measurement matches the local modeled measurement.

3) $\sigma_{il} \ge \sigma_{\text{ambig}} \ \forall \ l \ne j$; i.e., no other local modeled measurement matches the remote measurement.

Figure 10 includes conceptual illustrations of four cases, which (from left to right) show all conditions satisfied and violations of conditions 1, 2, and 3, respectively. Together, the conditions ensure that measurements are only assigned when observed and modeled measurements uniquely agree (with a statistical certainty determined by $\sigma_{\text{ambig}} > \sigma_{\text{match}} > 0$). StarFOX applied $\sigma_{\text{assign}} = 3$, $\sigma_{\text{safe}} = 6$, $\sigma_{\text{remove}} = 10$, $\sigma_{\text{match}} = 3$, and $\sigma_{\text{ambig}} = 6$.

Similar identity checks are needed to enable filter maneuver updates. ID numbers are assigned to targets by matching ground-provided orbit estimates to onboard estimates using Mahalanobis distance thresholds. SOD may then use the ID numbers associated with a priori maneuver plans to model maneuvers within the swarm.

Finally, to aid error recovery in degraded conditions, the autonomous health checks for the current state estimate are performed after each filter update. The filter attempts to automatically reinitialize if any of the following conditions are fulfilled:

1) The Mahalanobis distance between the most recent BOD estimate and current SOD estimate is above a user-specified threshold σ_{div}. TStarFOX applies $10 < \sigma_{\text{div}} < 100$. The BOD and SOD estimates are only compared if the time interval between them is sufficiently small (e.g., less than 1 orbit). This ensures consistency of BOD and SOD outputs.

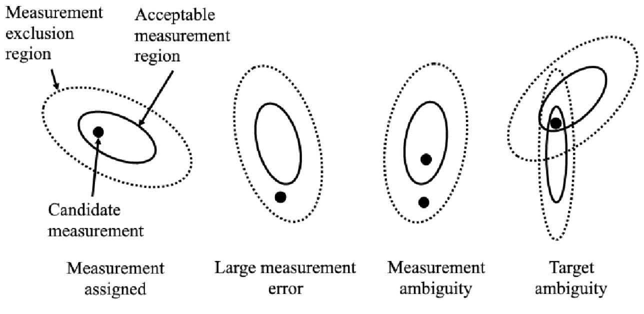

Fig. 10 Conditions in which all measurement assignment criteria are satisfied (leftmost) and conditions that violate each criteria (right).

2) The number of successive filter update steps that did not feature a measurement update is above a user-specified threshold. This ensures that ARTMS has not lost sight of its targets.

3) The number of successive filter measurement update steps that featured a large prefit residual is above a user-specified threshold. This ensures that ARTMS has not encountered tracking errors such that its state estimates are no longer reliable.

4) The SOD estimate is outside a user-specified valid range to ensure that it remains physically reasonable.

To re-initialize, SOD may either use ground-provided orbit estimates or the most recent BOD estimate.

IV. On-Orbit Tested

StarFOX is a core payload of the NASA Starling swarm mission, which consists of four propulsive 6U CubeSats. Starling was initially proposed in 2018 [19] as a testbed for autonomous swarming technologies, with four experiment payloads focused on decision-making, networking, maneuvering, and navigation. The swarm was launched into approximately sun-synchronous LEO on July 17, 2023, with an initial local time of the ascending node (LTAN) of 01:45. Nominal experiment operations commenced in November 2023 and continued until May 2024 [20], with a final LTAN of 06:00.

A. Bus Hardware

For StarFOX, each CubeSat carries two Blue Canyon Technologies Nano Star Trackers [8] (NST) aligned in antiparallel directions, either of which may collect imagery. NST parameters are provided in Table 2. During each experiment, one NST is chosen as the experiment camera and is pointed in the velocity or antivelocity direction to image swarm members. Images are provided to ARTMS with a nominal cadence of 60–120 s. The other NST provides a backup attitude solution.

StarFOX also receives inputs from the onboard GPS receiver, which optionally provides PVT solutions, and from the ISL, through CesiumAstro *S*-band software-defined radios. If the ISL is active, each spacecraft transmits a message to all other swarm members every 60 s and listens for messages from all other members. The payload processor for StarFOX is a Xiphos Q7S running at approximately 700 MHz. Finally, each satellite also possesses a cold gas propulsion system with four separate nozzles, used to conduct swarm phasing and reconfiguration maneuvers.

Figure 11 shows the four Starling spacecraft during integration at the NASA Ames Research Center. The spacecraft are referred to as SV1, SV2, SV3, and SV4 or, more affectionately, as Blinky, Pinky, Inky, and Clyde.

B. Swarm Geometry

Swarm geometry was determined by the differing needs of flight experiments and safety considerations [20]. To provide varying geometry for StarFOX, two distinct formations were employed: an in-train (IT) formation and a passive safety ellipse (PSE) formation. The IT formation features relatively little relative motion, with the swarm primarily separated in the velocity direction. The PSE formation provides additional relative motion (to aid angles-only observability) as well as passive safety via relative eccentricity/inclination vector separation [50]. The swarm was initially placed into the IT formation and then transitioned into the PSE formation in March 2024.

Table 2 Intrinsic parameters of the NST

Parameter	Value
Image size, pixels	1280×1024
FOV, deg	12×10
Pixel size, μm	5.3
Focal length, mm	30
Pixel space	10-bit, grayscale

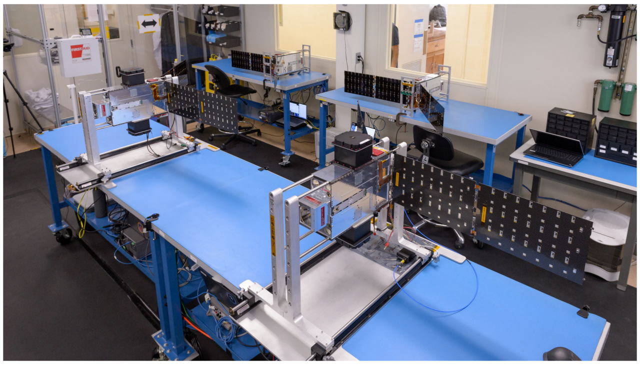

Fig. 11 The four Starling CubeSats during integration. Credit: NASA/ Dominic Hart (2022).

Table 3 Quasi-nonsingular OEs/ROEs of the swarm

Satellite	In-train formation, 02/05/24 00:00:00 UTC
SV4	$\boldsymbol{\alpha} = [6944\ \text{km}, -0.00004, 0.0016, 99.4\ \text{deg}, -151.1\ \text{deg}, -47.9\ \text{deg}]$
SV2	$\delta\boldsymbol{\alpha} = [21, -124350, 110, 202, 79, 1005]$ m
SV3	$\delta\boldsymbol{\alpha} = [-1, -79328, 42, 452, 36, 827]$ m
Satellite	Passive safety ellipse, on 05/08/24 at 00:00:00 UTC
SV4	$\boldsymbol{\alpha} = [6948\ \text{km}, -0.00009, 0.0013, 99.4\ \text{deg}, -38.0\ \text{deg}, -153.0\ \text{deg}]$ m
SV2	$\delta\boldsymbol{\alpha} = [23, -52266, -471, -523, 35, 1811]$ m
SV3	$\delta\boldsymbol{\alpha} = [-30, 155320, -297, -315, 16, 1144]$ m

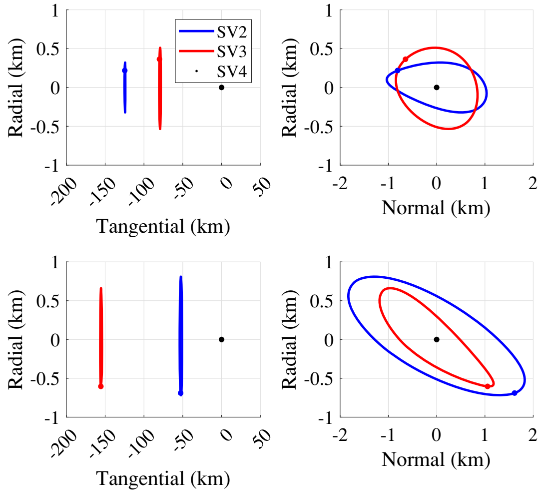

Fig. 12 Swarm relative motion in the IT formation (upper) and PSE formation (lower). ©2025 IEEE.

Table 3 presents OEs for SV4 and ROEs for SV2 and SV3 (with respect to SV4), following the conventions of Eqs. (3) and (4). SV1 was not involved in StarFOX operations due to on-orbit hardware complications. Figure 12 displays the relative motion geometry of the swarm in the RT and RN planes of SV4 in curvilinear coordinates.

C. Experiment Planning

In its current form, StarFOX proposes four primary navigation modes, as enabled by ARTMS:

1) Single observer, partially autonomous: Initial target relative orbit estimates are provided by the ground and refined on board by IMP + SOD. GPS is used for observer absolute orbit estimation.

2) Single observer, fully autonomous: Initial target relative orbit estimates are generated on board by IMP + BOD and refined on board by IMP + SOD. GPS is used for observer absolute orbit estimation.

3) Multi-observer, partially autonomous: Multiple observers share measurements over the ISL to estimate target relative orbits with improved accuracy. GPS is used for observer absolute orbit estimation.

4) Multi-observer, fully autonomous: Multiple observers share measurements over the ISL to estimate target relative orbits with improved accuracy. Inter-satellite bearing angles are used for observer absolute orbit estimation.

Modes 1 and 3 are most applicable to cooperative swarm scenarios in Earth orbit; mode 2 is most useful for SSA in which unknown targets must be detected and tracked; and mode 4 is most useful for deep space applications in which GPS is unavailable. Fourteen experiment blocks were originally planned [22], differentiated by swarm geometry, visible targets, presence of maneuvers, presence of the ISL, presence of GPS updates, and presence of a priori relative orbit knowledge. In total, 18 experiment blocks were executed over 34 experiment days. Experiment complexity increased throughout the schedule, with early experiments focusing on single-observer ground-assisted navigation and later experiments focusing on multi-observer autonomous navigation. Each swarm member could also be commanded independently, to allow multiple simultaneous experiments to be executed.

D. Experiment Execution

Preparation for each experiment involved the fusion of data from three sources: NASA Flight Dynamics, NASA mission operators, and Stanford SLAB. Flight Dynamics produced predicted swarm ephemerides and planned maneuvers, which were used to provide each spacecraft with its absolute orbit estimate for initializing navigation and, optionally, a priori relative orbit estimates for its targets. NASA operators created command sequence plans in order to schedule StarFOX activities by the spacecraft bus, including the attitude profile, imaging cadence, crosslink radio activation, and ground station contacts. SLAB provided telecommand (TC) tables for ARTMS to govern on-orbit behavior, including activation of individual ARTMS algorithms and granular tuning of IMP, BOD, and SOD parameters. Subsequently, NASA's Mission Operations Systems team performed command validation using a hardware-in-the-loop simulation environment. Validated products were uplinked to the swarm.

StarFOX payload operations during each experiment encompassed several tasks. First, spacecraft were slewed to align their star trackers with the (anti)velocity direction to ensure consistent swarm visibility. Command products were then unpacked and verified in order to commence nominal operations, e.g., taking star tracker images, running ARTMS applications, and saving outputs as telemetry (TM). In practice, StarFOX operations could be interrupted by various tasks onboard, such as swarm maneuvers and ground passes (affecting satellite attitude) or Globalstar signaling (which reduced GPS accuracy). Experiments continued for ~24 h until a daily payload reset. Onboard TM was downlinked during ground passes and provided to Stanford SLAB for analysis.

Table 4 illustrates a typical operational cadence for ARTMS, with order-of-magnitude runtimes. Each cycle begins with acquisition of images from both star trackers, followed by IMP execution, ISL broadcast, and SOD execution, once per minute, at a set offset to allow the previous step to complete. BOD runs in the background once every 90 min. Onboard computational loads were sufficiently low to enable feasible real-time operation.

Table 4 Example timing for onboard ARTMS operations

Parameter	Sample time, s	Runtime, s	Offset, s
Clock input	1	— —	— —
GNSS input	1	— —	— —
VBS image 1	60	— —	0
VBS image 2	60	— —	+10
IMP execution	60	10	+20
ISL broadcast	60		+30
SOD execution	60	10	+40
BOD execution	5400	600	Background

V. On-Orbit Performance

StarFOX flight results are generated in two ways. In the first case, ARTMS produces telemetry while operating in orbit, and telemetry is downlinked and analyzed on the ground. Navigation performance is numerically assessed by comparing the ARTMS orbit estimate in telemetry to "ground truth" orbit states, produced by the NASA Flight Dynamics team via batch processing of GPS flight data. Expected position errors and uncertainties in the ground truth are significantly less than 10 m (1σ).

The second case is a digital twin of the swarm, which operates as similarly as possible to ARTMS in flight. ARTMS inputs (i.e., star tracker images and metadata, GPS data, timing data, and ISL data) are downlinked and are processed in a flight-like manner using ARTMS flight software running on the ground. The ARTMS twin operates in a MATLAB/Simulink environment on a PC, which interfaces with the C++ flight code via s-functions. Four separate instances of ARTMS can be simulated in parallel to model the complete swarm.

Table 5 presents the dynamics models applied within IMP, BOD, and SOD during flight experiments. Adaptive process noise estimation was not applied during these initial tests.

The performance of ARTMS is characterized across a quartet of representative scenarios: a single-observer, multi-target experiment in navigation mode 1; a single-observer, single-target experiment in navigation mode 2; a multi-observer, multi-target experiment in navigation mode 3; and a multi-observer, multi-target experiment in navigation mode 4. Outcomes provide the first-ever demonstrations of optical angles-only navigation for a satellite swarm in orbit.

A. Relative Navigation: Single-Observer, Multi-Target

Figure 13 presents in-flight estimation performance from SV4 telemetry on 05/16/2024. SV4 points its star tracker in the antivelocity direction to image SV2 (at 50 km ISD) and SV3 (at 130 km ISD). Relative orbit estimates for SV2 and SV3 are initialized using stale information from the ground (approximately 3–4 days old). GPS is used to maintain SV4's absolute orbit estimate, and relative navigation is performed using bearing angles from SV4 only. Images were taken every 60 s.

It can be seen in Fig. 13 that SV4 successfully tracks and performs orbit estimation for multiple targets simultaneously. It is also able to maintain a converged relative orbit estimate for 21 h without requiring maneuvers. This is the first time these capabilities have been demonstrated in flight.

Performance displays expected trends for angles-only navigation. Error is primarily captured by the $\delta\lambda$ ROEs, which corresponds to the weakly observable target range. Uncertainties (1σ) for $a\delta\lambda$ at the end of the experiment period are 1.3% for SV2 and 2.6% for SV3, as a percentage of the target range. Uncertainties in other ROEs are less than 0.02% of the target range because these components are more strongly observable. Estimation errors generally remain within 1σ bounds, indicating good filter health. Errors and uncertainties are also somewhat proportional to range in practice, because at longer distances, the same sensor noise floor corresponds to a larger geometric uncertainty.

Table 5 Dynamics models for ARTMS in flight

Model	Perturbations	Propagation
IMP	Keplerian, none	Analytic
BOD	J_2 gravity [34]	Analytic
SOD	10×10 GGM01S gravity [2] Harris–Priester atmosphere with cannonball drag, $B_D \approx 0.015$	RK4 integrator 30 s time step

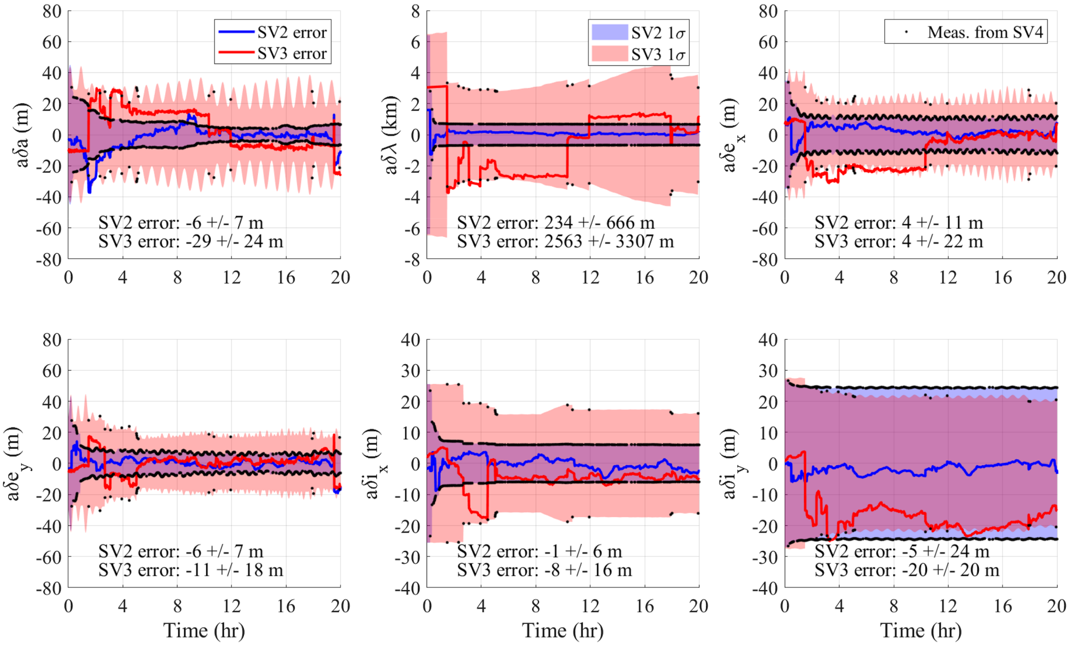

Fig. 13 ROE estimation errors and uncertainties (1σ) for the single-observer, multi-target case. ©2025 IEEE.

Figure 14 presents pre- and postfit residuals for the SV2 target. Azimuth residuals are generally less than 50”, whereas elevation residuals are larger, which suggests that additional uncertainties or error sources are affecting the elevation angle in the VBS frame. However, no obvious biases are present. In this experiment, bearing angle measurement noise of $40''$ (1σ) was assumed in the filter, because the prior value of $20''$ [8,22] from preflight testing was found to be insufficiently robust.

Results were partially affected by poor target visibility, most obviously for SV3. The visual magnitude of the swarm was significantly dimmer than expected. As a result, targets with an ISD greater than 100 km were not reliably detected by IMP. Figure 15 illustrates this phenomenon: SV3 is consistently visible above background noise, whereas the more distant SV2 is not. In Fig. 13, fewer than 20 SV3 measurements were obtained over 12 orbits, as indicated by black dots. State uncertainties for SV3 are therefore larger, and there are step changes in the SV3 estimation error that correspond to individual sparse measurements. Nevertheless, ARTMS was able to maintain a reasonable orbit estimate for both targets, validating the robustness of the architecture and its ability to maximize angles-only observability. A direct comparison between in-flight performance and a simulation featuring more idealized measurement conditions is presented in [24]. Estimation performance for SV2 remains similar in flight and in simulation, whereas performance for SV3 is vastly improved in simulation (achieving steady-state $\delta\lambda$ errors of 652 ± 1083 m) due to more consistent measurement availability.

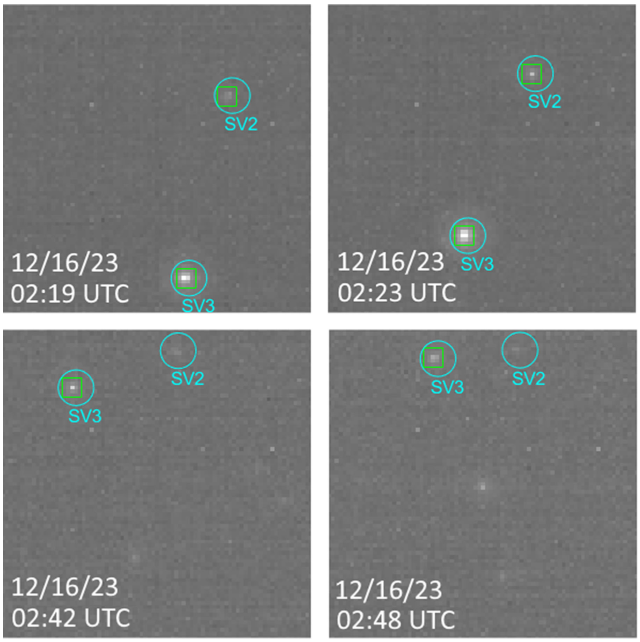

Fig. 15 Target visibility in flight images from SV4. Green boxes indicate successful detection by IMP.

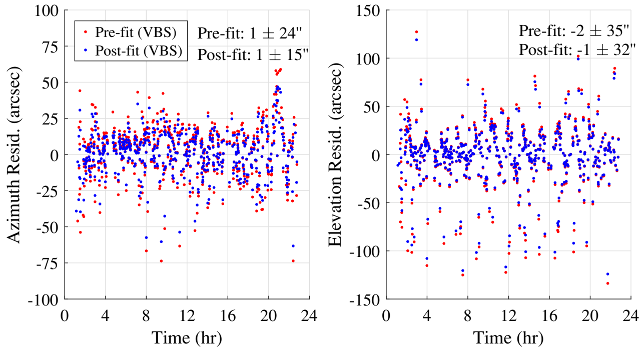

Fig. 14 Pre- and postfit measurement residuals for the SV2 target in the single-observer, multi-target case.

B. Relative Navigation: Autonomous Initialization

Autonomous initialization of relative navigation for an unknown target was achieved on 05/08/24, which is the first time this capability has been demonstrated for an angles-only target in flight. SV2 points its star tracker in the velocity direction to image SV4 (at 50 km ISD). No a priori relative orbit information for SV4 is provided. SV2 performs multi-hypothesis tracking with IMP to detect potential targets in the FOV. Every 90 min, BOD attempts to compute a relative orbit initialization using IMP measurement batches and the PVT solution from the GPS receiver. If the computed orbit is within state and uncertainty tolerances, it is used to initialize SOD, and formal navigation commences. GPS is subsequently used to maintain SV2's absolute orbit estimate. Images were taken every 60 s.

Figure 16 presents overlaid regions of interest from one orbit of images, as recorded in IMP telemetry. The point sources corresponding to SV4's relative orbit are bordered by the ellipse. Variations in intensity correspond to variations in spacecraft attitude and solar phase angle during the orbit. The relative motion ellipse is clearly visible, despite the rate of missing measurements on board being higher than anticipated. Nominally, 1320 images would be processed by IMP throughout the 22 h experiment, whereas in practice, only 563 images were processed during the same timespan on 05/08/24, implying a success rate of 43%. Due to unexpectedly high processing loads in the Starling flight software, star tracker images were not always successfully transferred to the StarFOX payload. The target was subsequently observed in only 338 images, reducing measurement rates to 26% of the nominal rate. Image quality was additionally degraded by low signal-to-noise ratios and the presence of many hot pixels in images. In the worst case, nearly 0.5% of all pixels in the image were affected, and it became necessary to update IMP to filter hot pixels more effectively.

Despite these challenges, IMP successfully performed MHT by applying SAMUS to unidentified bearing angles. Figure 17 presents track hypotheses for a single visible target over the first 8 h, where distinct colors indicate distinct valid hypotheses and distinct black markers indicate distinct pruned hypotheses. The blue, red, and green segments follow the expected elliptical kinematics of a target in a similar orbit and could be connected into a single track. In contrast, the black segments display large velocity or direction changes and are considered physically unreasonable. Of note is the track made up of repeated measurements near $[0,0]^\top$, which persisted as a possible in-train target (i.e., a target with very little relative motion) for 36 min before being pruned.

IMP performance is assessed by characterizing bearing angle measurement errors in $\mathcal{I}$ compared to ground truth bearing angles in $\mathcal{I}$ derived from GPS, as in Fig. 18. Some systematic errors are present in azimuth and elevation, partially due to a drift in image time-tags; the star tracker time did not stay perfectly synchronized with the onboard clock (which itself was synchronized to GPS time), meaning that image time-tags became inaccurate. This drift was observed to vary between experiments and was approximately +300 ms on 05/08/24. A separate software update was necessary to allow user-specified corrections to the time drift on board. Measurement errors from the digital twin after correcting for this drift are shown on the right-hand plot in Fig. 18. The bias is reduced but remains nonzero, potentially due to changes to star tracker calibration postlaunch.

BOD performance is summarized in Table 6, which displays the target $\delta\lambda$ estimates produced by BOD. The first successful initialization occurred 7.75 h after the experiment started, at the third recorded BOD call. (Some intermediate BOD calls were not recorded due to insufficient measurements or unavailable telemetry.) Inconsistent observation rates in IMP meant correct tracks were often pruned too early and were only rarely propagated above the minimum length of 20 measurements required for BOD to run. Despite these obstacles, BOD was able to produce a first valid initialization using only 28 measurements, with $a\delta\lambda$ errors of less than 20% of the 53 km ISD.

BOD outputs were also impacted by poorer-than-expected GPS performance. During 05/08/24, mean position uncertainties reported by the SV2 GPS receiver were 9.79 ± 3.61 m/s and mean velocity uncertainties were 1.16 ± 0.41 m/s (1σ). Though position uncertainties were in line with preflight expectations [22], velocity uncertainties were two orders of magnitude larger, corresponding to large uncertainties in the semimajor axis. As opposed to SOD, which

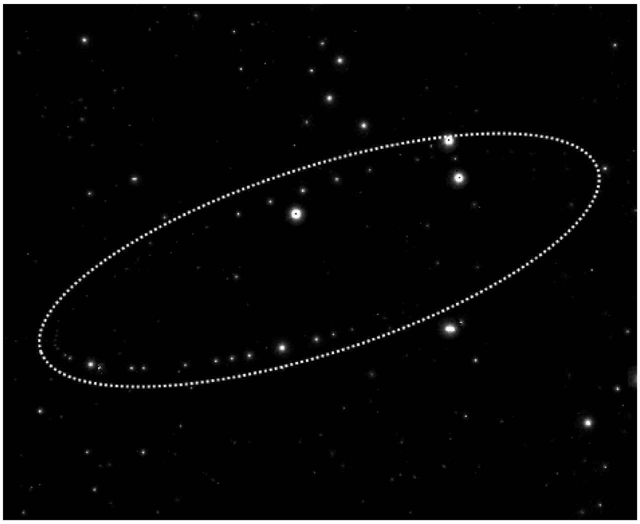

Fig. 16 Overlaid regions of interest from one orbit of images, for SV2 viewing SV4 in the PSE formation.

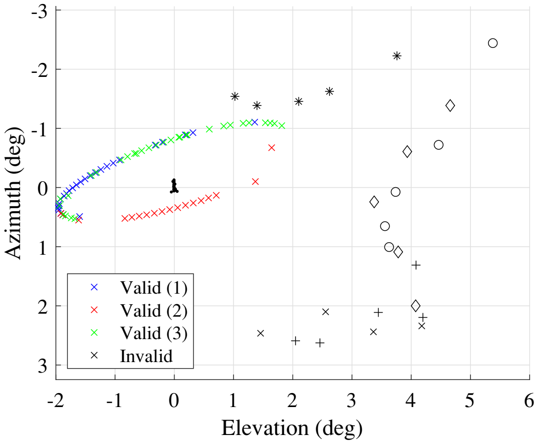

Fig. 17 Track hypotheses produced by SAMUS. ©2025 IEEE.

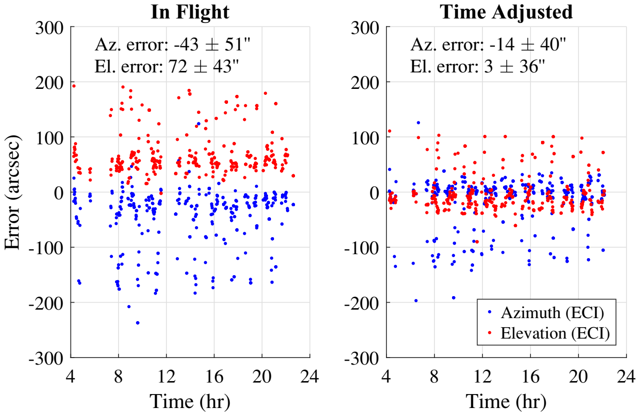

Fig. 18 Measurement errors for target bearing angles in the ECI frame produced by IMP.

Table 6 BOD estimation performance for $\delta\lambda$

Time, h	Meas.	Span, h	$a\delta\lambda$ error, km	$a\delta\lambda$ uncert. (1σ), km
1.75	24	0.9	25.45	56.26
6.25	22	2.0	−14.16	15.66
7.75	28	3.5	−8.33	9.64
9.25	42	1.5	−20.51	24.23
10.75	67	3.0	−5.92	11.53
12.25	105	4.5	−6.14	7.53
13.75	135	6.0	−6.60	5.80
16.75	145	7.5	−7.03	4.72
19.75	159	7.8	−4.44	3.94

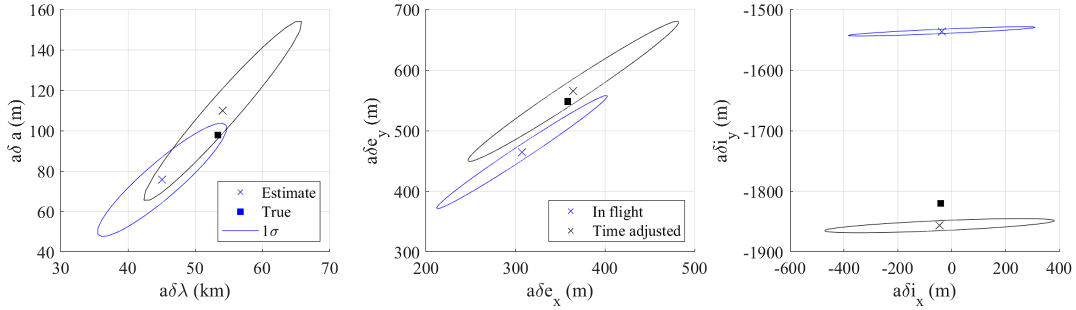

Fig. 19 ROE estimates and 1σ uncertainties from the BOD initialization.

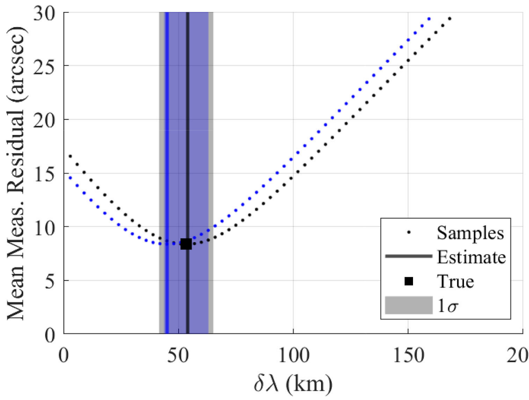

Fig. 20 Measurement residuals for $\delta\lambda$ samples in BOD for in-flight (blue) and time-adjusted (black) results.

sequentially filters measurements to reduce velocity uncertainties over time, BOD relies on only the current PVT solution from the observer's GNSS receiver when computing an initialization. Uncertainty in the absolute orbit estimate becomes coupled to uncertainties in the BOD relative orbit estimate, and as a result, output BOD uncertainties were often too large to be acceptable for initializing SOD. To compensate, it was necessary to tune BOD to reduce output state uncertainties by a factor of 5, while allowing BOD to refine its semimajor axis estimate as part of the initialization process. Subsequently, errors were generally well-characterized by the estimated 1σ uncertainty. BOD is therefore relatively robust to large errors in a; an error of 1 m/s in velocity corresponds to a shift of more than 1 km in the semimajor axis at Starling's altitude. An exception is when measurement batches span larger time periods, as in the last three rows of Table 6, whereupon BOD can become overconfident.

Estimation performance for the initialization at $t = 7.75$ h is visualized in Figs. 19 and 20. Figure 19 displays true ROEs, estimated ROEs, and estimated 1σ bounds for flight telemetry (blue) and the on-ground twin with corrected time tags (black). Correcting for the star tracker time drift significantly reduces errors in the estimated ROEs. After adjustment, the true ROE value is within 1σ of the estimated ROE value for $[\delta a, \delta\lambda, \delta e_x, \delta e_y, \delta i_x]$. An exception is the δi_y value, which lies outside the 3σ uncertainty bound. This suggests that additional on-orbit error sources remain unaccounted for when compared to the robust BOD results obtained during preflight simulations [22]. Figure 20 displays the mean residuals corresponding to individual $\delta\lambda$ samples, mirroring the color scheme of Fig. 19. The expected convex behavior is evident, and selecting the sample that produced the least residuals led to a reasonable initialization for the weakly observable $\delta\lambda$ term.

SOD performance postinitialization is visualized in Fig. 21. The blue line corresponds to estimation errors from in-flight telemetry. It is clear that even if state uncertainties are converging, a strong bias is present in the state error for $\delta\lambda$, δe_x, δe_y, and δi_y due to image timing errors on board. The black line corresponds to the on-ground twin after timing errors are corrected, which removes the bias in all relevant ROE estimates. Target measurements in flight are indicated by vertical gray bars, indicating lengthy periods without new measurements.

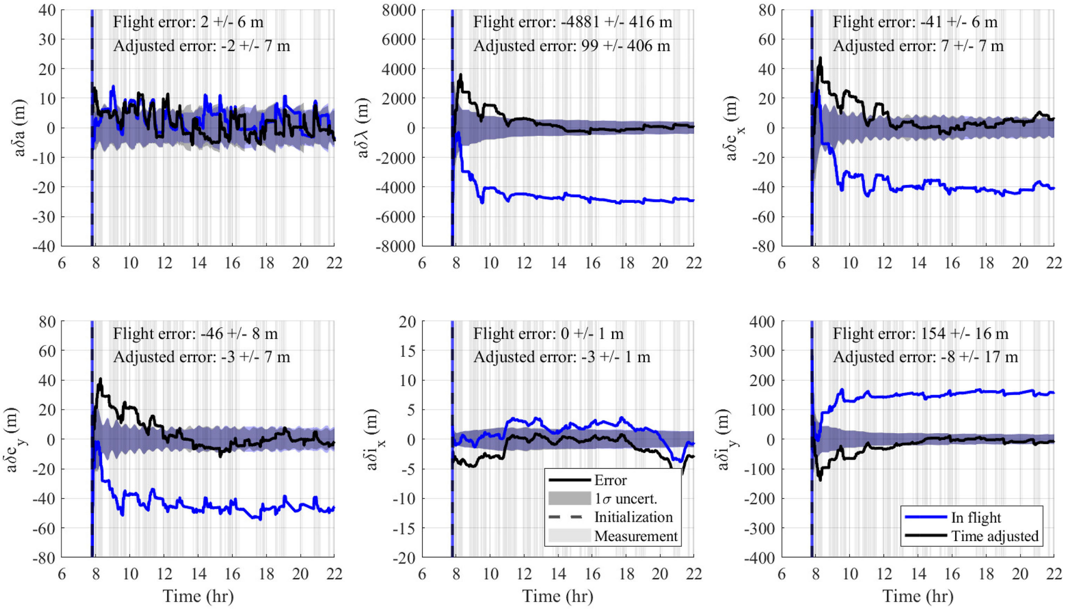

Fig. 21 ROE estimation errors and uncertainties (1σ) for the autonomous initialization case.

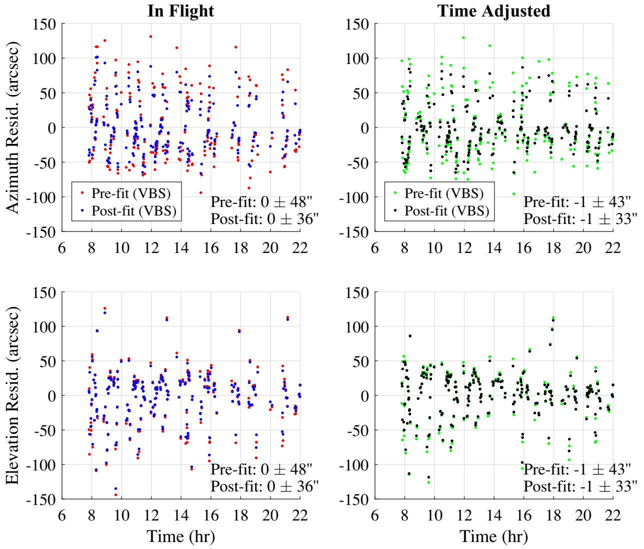

Fig. 22 SOD pre- and postfit measurement residuals for the autonomous initialization case.

Filter measurement residuals are plotted in Fig. 22. Residuals appear to be zero-mean, and postfit residuals are consistently smaller than prefit residuals. The image timing errors do not produce an obvious bias in the residuals themselves, suggesting that the time drift is unobservable with the available inputs. This is consistent with previous work by the authors [30,31], which has shown that differential clock offsets between observers are observable using inter-satellite bearing angles, whereas absolute clock offsets are unobservable with inter-satellite bearing angles.

C. Relative Navigation: Multi-Observer, Multi-Target

Figure 23 presents estimation performance from SV4 telemetry on 02/05/2024. In this scenario, SV4 points its star tracker in the antivelocity direction to image SV2 (at 130 km ISD) and SV3 (at 80 km ISD). Relative orbit estimates for SV2 and SV3 are initialized using stale information from the ground (3–4 days old). Relative orbit estimation is performed using bearing angles from multiple observers over the ISL, such that SV4 additionally receives ISL measurements from SV2 (measuring SV3 and SV4) and SV3 (measuring SV2). GPS is used to maintain SV4's absolute orbit estimate, and images were taken every 120 s.

The colored dots in Fig. 23 mark instances when a measurement broadcast by a remote observer was used to update local state estimates on SV4. It can be seen that SV4 utilizes measurements from all three spacecraft, which is the first time this distributed and adaptive multi-observer capability has been demonstrated in flight. Results reveal certain advantages of multi-observer navigation. Firstly, SV4 only obtains five local measurements of SV2 throughout the experiment period. It must instead rely on measurements of SV2 from SV3 (and SV3 from SV2) to update its onboard orbit estimate for SV2. Thus, multiple observers are able to provide added robustness when measurements are sparse (compared to, e.g., SV3 in Fig. 13). Secondly, state uncertainty is reduced compared to an equivalent single-observer case. Uncertainties (1σ) for $a\delta\lambda$ at the end of the experiment period are 0.6% for SV2 and SV3, as a percentage of the target range. Uncertainties in other ROEs are less than 0.01% of the target range. The additional geometric information provided by multi-observer measurements significantly improves angles-only observability.

Anomalies are also evident in Fig. 23, such as sudden decreases in the $a\delta\lambda$ uncertainty at several time instants, e.g., at $t = 8$ h. This is due to unintended behavior when the SOD filter resets. As part of the ISL message, swarm members broadcast their absolute orbit estimate, which allows the receiving spacecraft to match any remote measurements to local state estimates. Typically, this orbit estimate is retrieved from the SOD filter. However, when SOD re-initializes (e.g., due to a large number of skipped measurement updates), the SOD estimate becomes temporarily invalid, and the onboard GNSS solution is sent instead as backup. When SV4 received this GNSS solution on 02/05/24, it was interpreted as a measurement that could be used to update the relevant target's state estimate. GNSS improves observability in this scenario, which led to sudden decreases in target state uncertainty. On-ground tests with the digital twin, which explicitly disable remote GNSS updates, do not display the same step changes and indicate that, after 20 h, the steady-state angles-only $\delta\lambda$ error is 148 ± 2052 m for SV2 and 56 ± 768 m for SV3 [24].

The initial $\delta\lambda$ uncertainty for SV2 is also large. As summarized in Sec. IV.D, the initial relative orbit estimates for this scenario were computed using a priori orbit and maneuver information uplinked from the ground several days prior. Errors in the a priori data and onboard maneuver execution were larger than anticipated, producing initial state errors outside typical uncertainty bounds.

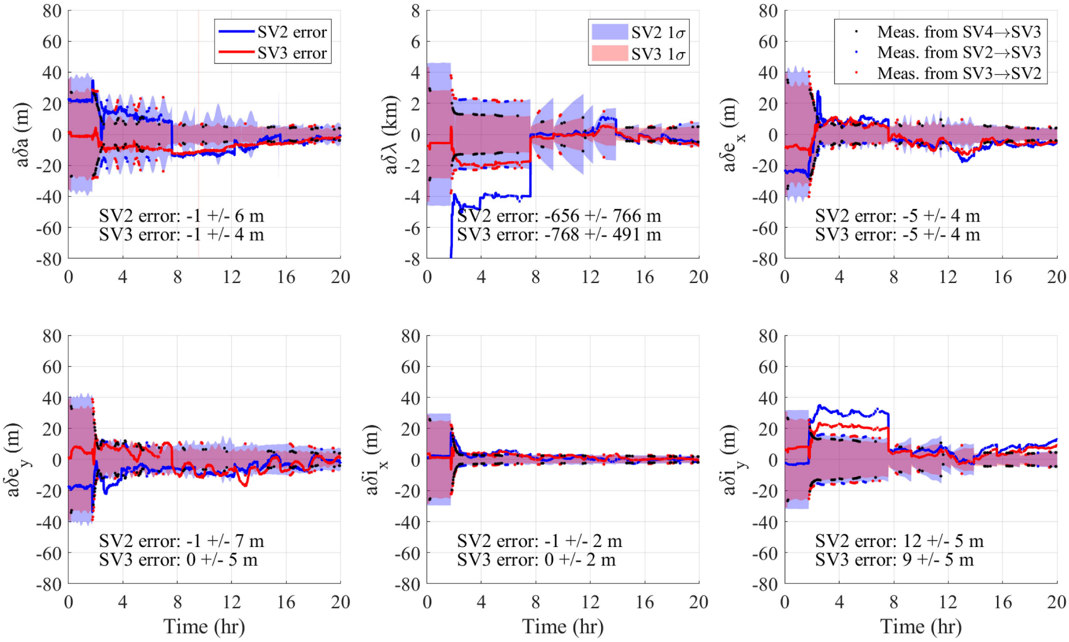

Fig. 23 ROE estimation errors and uncertainties (1σ) for the multi-observer, multi-target case. ©2025 IEEE.

D. Absolute Navigation: Multi-Observer, Multi-Target

Demonstrations of navigation mode 4 were attempted on 05/09/24 and 05/16/24 in flight, but crosslink issues prevented successful execution. It was later demonstrated using flight data from 03/24/24, postprocessed by the ARTMS digital twin running on the ground. The following scenario applies SV2, SV3, and SV4 to process approximately 3300 star tracker images over 40 h, with simulated ISL transmissions between all members. Onboard state estimates are initialized using artificially degraded orbit knowledge to emulate a coarse initialization with uncertainties on the scale of several kilometers. After initialization, no additional external absolute orbit knowledge is provided (i.e., no GPS data is processed). Both absolute and relative orbit determination are performed using bearing angles only.

Figure 24 presents absolute orbit estimation performance for SV4. Orbit element errors are normalized by the observer’s semimajor axis for geometric interpretation. Absolute orbit estimation using only inter-satellite bearing angles is particularly challenging due to the weak observability of the problem. However, if multiple observers are present and are obtaining measurements of a common target (and each other), enough geometric information is present to facilitate long-term absolute orbit convergence without external updates [11,22,51].

Here, SV4 observes SV2 and SV3 as targets but also receives crosslink measurements from SV3 (observing SV2 and SV4) and SV2 (observing SV3). The coarse state initialization is successfully refined over the experiment period by leveraging multi-observer images. In the RTN frame of SV4, this corresponds to an initial position error of 6950 m, which is reduced to a final error of 1200 m. Repeating the simulation with halved initial errors produces final position errors of 850 m. Convergence of the state estimate is slow due to challenging observability, but estimation errors remain within 1σ uncertainty bounds, indicating reasonable filter health. This is the first time this capability has been demonstrated using a flight software digital twin and flight data.

Figure 25 displays the ISL measurements received by SV4. The most consistently visible link was SV4 observing SV2, whereas other observations were more sparse due to the visibility issues previously described. The variation in measurement rates between seemingly equivalent spacecraft pairs (e.g., SV4 → SV3 and SV3 → SV4) illustrates an additional difficulty when applying optical measurements: visibility varies significantly with orbit geometry, measurement geometry, sensor attitude, and parameters of the specific VBS at hand.

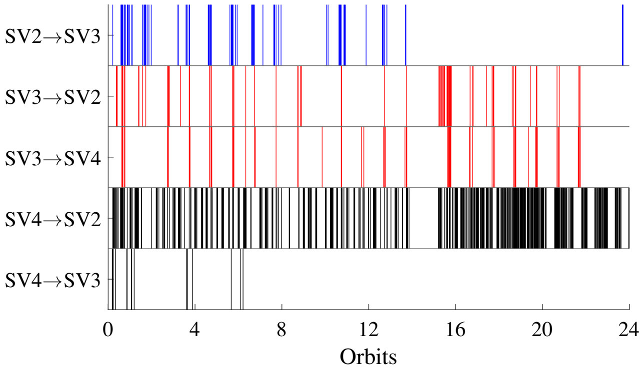

Fig. 25 Measurements received by SV4. A bar indicates that a measurement was received at that epoch.

E. Discussion

Eleven StarFOX experiment goals were defined before launch, informed by Starling objectives and preflight performance expectations [22]. These goals, with respect to StarFOX outcomes, are summarized in Table 7. All thresholds were successfully achieved, and the majority of goals were successfully achieved through flight experiments or postprocessing of flight data [23,24].

Objectives 6 and 7, which relate to swarm maneuvers, have not yet been discussed. Station-keeping maneuvers were executed during StarFOX experiments on 05/09/24 and 05/16/24, with the objective of observing a corresponding decrease in state uncertainty. Known maneuvers are expected to improve angles-only observability by aiding the filter in disambiguating target range [9,15]. Although maneuvers were successfully tracked on board, the executed delta-v’s were small ($\sim$0.01 m/s) and did not significantly improve state uncertainties as a result.

StarFOX has encompassed the first on-orbit demonstrations of 1) convergence of angles-only state estimates without maneuvers, 2) angles-only navigation for multi-target and multi-observer systems, and 3) initialization of angles-only relative navigation for an unknown target. The first demonstrations of angles-only absolute orbit maintenance were also achieved via a digital twin operating on flight data. This acts to greatly increase the technology readiness level of angles-only navigation for spacecraft. Furthermore, the overall autonomy, flexibility, and accuracy demonstrated by ARTMS are strong enablers for future distributed missions. This is especially relevant for missions operating autonomously in deep

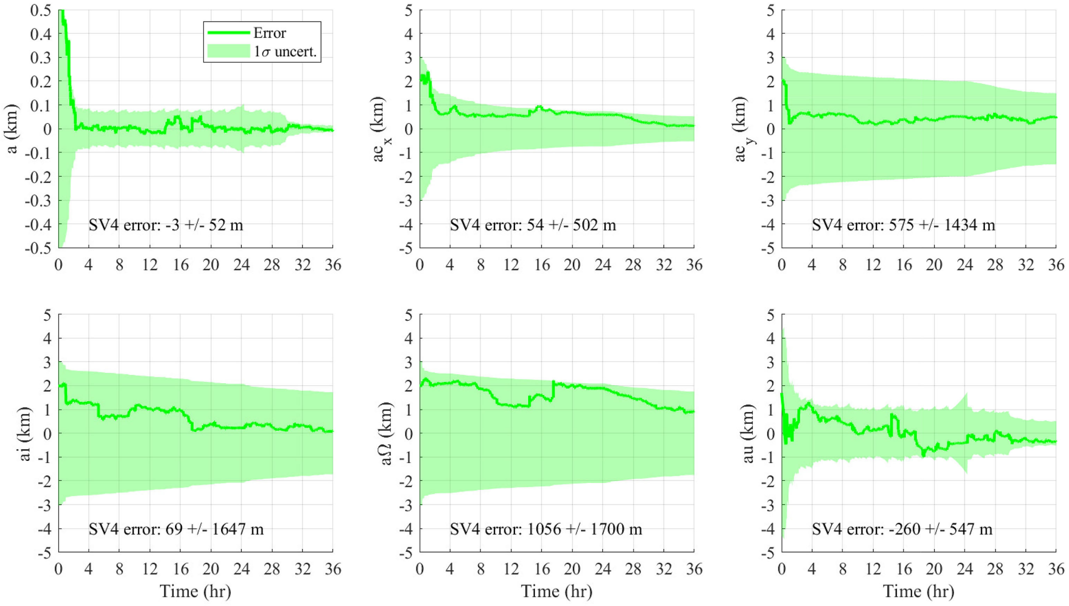

Fig. 24 OE state estimation errors and uncertainties (1σ) for the multi-observer case.

Table 7 Summary of StarFOX performance goals and flight results (©2025 IEEE)

No.	Objective	Threshold	Goal	Result
1	Tracking multiple targets simultaneously	1 target	3 targets	2 targets (in flight)
2	Accuracy of relative position knowledge with one observer	Estimate produced	1% error relative to ISD	0.5% error relative to ISD (in flight)
3	Convergence time for relative position knowledge with one observer	Convergence observed	6 orbits	2 orbits (in flight)
4	Accuracy of relative position knowledge with multiple observers	Estimate produced	0.1% error relative to ISD	0.5% (in flight), 0.1% (on ground)
5	Convergence time for relative position knowledge with multiple observers	Convergence observed	2 orbits	1 orbit (in flight)
6	Accuracy of relative position knowledge in presence of maneuvers	Estimate produced	0.1% error relative to ISD	Estimate produced (in flight)
7	Convergence time for relative position knowledge in presence of maneuvers	Convergence observed	2 orbits	Convergence observed (in flight)
8	Accuracy of absolute position knowledge with multiple observers	Estimate produced	1 km error	850 m error (on ground)
9	Convergence time for absolute position knowledge with multiple observers	Convergence observed	24 h	Convergence observed (on ground)
10	Time to produce autonomous relative orbit initialization	Initialization produced	2 orbits	5 orbits (in flight)
11	Accuracy of target range estimate for relative orbit initialization	Initialization produced	20% error relative to ISD	16% (in flight), 2% (on ground)

space (i.e., without access to external measurement sources) and missions focusing on space situational awareness, which aim to track passive objects. Relative position uncertainties of tens of meters in the radial/cross-track directions and hundreds of meters in the along-track direction are sufficient to enable a variety of mission objectives [30,31], and vision-based methods may also serve as a robust secondary navigation system.

StarFOX also provides a number of lessons learned from the challenges encountered in flight. In particular, when leveraging optical sensors, there must be a focus on incorporating physically realistic image simulation capabilities. Accurate modeling of target visibility and image noise is necessary during development to guarantee consistent measurement availability in flight. For StarFOX, the nominal ISD was too large to support effective navigation in many cases, which could not be mitigated within the constraints of the Starling concept of operations.

There must also be the capacity to stimulate the optical sensor during integration and testing. Although extensive attention was paid to these aspects when developing the flight software (including hardware-in-the-loop testing using a Blue Canyon NST and an optical stimulator testbed [22]), it could not be extended to the integration of the satellites themselves. As a result, sources of measurement error and timing errors on board the satellites were not well understood. This led to estimation errors in flight. Unexpected behaviors (such as unreliable image transfers) were also encountered due to incomplete characterization of computation loads on the payload avionics.

The success of the StarFOX despite these issues illustrates the importance of flexibility and robustness during system design. Efforts were made during the development of ARTMS to allow the system to adapt to differing conditions, either autonomously (via fallback execution paths, contingencies, and data validity checks) or in a user-defined fashion (via tunable telecommand parameters). These efforts enabled successful outcomes even when on-orbit conditions proved unfavorable. Furthermore, access to an on-ground digital twin became vital. The ability to receive telemetry, perform playback, and rapidly replicate on-orbit behavior provided in-depth understanding of the issues encountered, and software fixes could be deployed with swift turnarounds. This digital twin (in concert with high-fidelity simulation frameworks) has also facilitated characterizations of experiment performance when measurement conditions are closer to nominal [24]. The resulting outputs closely match the performance expected from preflight verification, validating the overall approach [22].

Finally, it should be noted that StarFOX employed CubeSat star trackers, radios, and avionics for navigation. This demonstrates that angles-only navigation is achievable using miniaturized hardware, but improved performance could likely be achieved if using a more sensitive and more accurate VBS. Similarly, the increased processing power afforded by a dedicated GPU or FPGA would allow faster image sample rates and processing of higher-resolution data.

VI. Conclusions

This paper presents an overview of the StarFOX, the first in-flight demonstration of autonomous angles-only navigation for a spacecraft swarm. StarFOX is a core experiment payload of the NASA Starling mission, which consists of four propulsive 6U CubeSats launched in July 2023. StarFOX flight data exhibits four critical advances over the prior state-of-the-art for angles-only navigation, including the first demonstrations of 1) multi-target and multi-observer navigation, 2) autonomous onboard initialization of relative navigation for an unknown target, 3) long-term maneuver-free convergence of orbit estimates, and 4) simultaneous absolute and relative orbit determination using angles-only measurements.

StarFOX applies ARTMS, which provides autonomous, distributed, scalable navigation for distributed systems in deep space using inexpensive optical sensors. In flight, ARTMS was successfully applied to achieve complete swarm orbit determination with minimal external orbit information and without maneuvers. Star trackers were used to image swarm members and obtain bearing angle measurements, which were optionally shared over the crosslink. Observers were able to navigate for multiple target spacecraft and displayed maneuver-free angles-only convergence, achieving relative position estimation uncertainties of 1.3% of target range (1σ). Multiple observers were able to cooperate to estimate relative orbits with improved 0.6% position uncertainty relative to range (1σ). Unknown targets were detected and tracked onboard, and bearing angles were used to autonomously compute a relative orbit initialization. Multi-observer bearing angles were leveraged to refine absolute orbit uncertainties simultaneously, from several kilometers initially down to less than 1 km over 48 h.

Critical challenges were also encountered in flight. Target visibility and signal-to-noise ratios were poorer than expected, leading to sparse measurement availability and degraded state accuracy. Image timestamps also became desynchronized with GPS time, producing a state bias if unaccounted for. Additional issues included unreliable image transfers due to high processing loads. These problems could likely have been mitigated if an optical stimulator had been available during integration and testing to allow verification of ARTMS running on satellite avionics with images in the loop. Furthermore, improved radiometric consistency when simulating images in software would have led to a more realistic characterization of visibility during the mission design stage. StarFOX’s successes in spite of these challenges were made possible by the robustness of its software design, including fallback logic in degraded conditions, and the use of a digital twin on the ground to play back flight data and reproduce on-orbit behavior. These lessons may inform the development of future deep space or space situational awareness missions intending to rely on optical navigation techniques. Overall, StarFOX represents a significant step toward on-orbit autonomy and the more widespread application of vision-based navigation in space.

Acknowledgments

This work was supported by the NASA Small Spacecraft Technology Program cooperative agreement number 80NSSC18M0058 and by the Air Force Office of Scientific Research award number FA9550-21-1-0414. The authors acknowledge Joshua Sullivan and Adam W. Koenig for their invaluable contributions to the development of ARTMS and StarFOX. They also thank Soon S. Hwang and the entire NASA Ames Research Center Starling team for their tireless support of StarFOX efforts.

References

[1] D’Amico, S., Pavone, M., Saraf, S., Alhussien, A., Al-Saud, T., Buchman, S., Byer, R., and Farhat, C., “Miniaturized Autonomous Distributed Space System for Future Science and Exploration,” *8th International Workshop on Satellite Constellations and Formation Flying*, International Astronautical Federation’s Astrodynamics Committee, Delft, The Netherlands, 2015.

[2] 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

[3] Krieger, G., Moreira, A., Fiedler, H., Hajnsek, I., Werner, M., Younis, M., and Zink, K., “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

[4] 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, No. 3, 2012, pp. 834–850. https://doi.org/10.2514/1.55638

[5] Burch, J., Moore, T., Torbert, R., and Giles, B., “Magnetospheric Multiscale 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

[6] Holzinger, M. J., and Jah, M. K., “Challenges and Potential in Space Domain Awareness,” *Journal of Guidance, Control, and Dynamics*, Vol. 41, No. 1, 2018, pp. 15–18. https://doi.org/10.2514/1.G003483

[7] Arney, D., Sutherland, R., Mulvaney, J., Steinkoenig, D., Stockdale, C., and Farley, M., “On-Orbit Servicing, Assembly, and Manufacturing (OSAM) State of Play,” NASA TR 20210022660, 2021.

[8] Palo, S., Stafford, G., and Hoskins, A., “An Agile Multi-Use Nano Star Camera for Constellation Applications,” *Proceedings of the Small Satellite Conference*, Paper SSC13-III-5, AIAA/USU, 2013.

[9] 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

[10] Sullivan, J., and D’Amico, S., “Adaptive Filtering for Maneuver-Free Angles-Only Navigation in Eccentric Orbits,” *27th AAS/AIAA Space Flight Mechanics Conference*, AAS Paper 17-402, Springfield, VA, 2017.

[11] Kruger, J., and D’Amico, S., “Observability Analysis and Optimization for Angles-Only Navigation of Distributed Space Systems,” *Advances in Space Research*, Vol. 73, No. 11, 2023, pp. 5464–5483. https://doi.org/10.1016/j.asr.2023.08.055

[12] Greaves, J., and Scheeres, D., “Spacecraft to Spacecraft Absolute Tracking for Autonomous Navigation of a Distributed Space System from Relative Sensors,” *Journal of the Astronautical Sciences*, Vol. 71, No. 46, Oct. 2024, p. 46. https://doi.org/10.1007/s40295-024-00463-6

[13] Koenig, A. W., and D’Amico, S., “Observability-Aware Numerical Algorithm for Angles-Only Initial Relative Orbit Determination,” *2020 AAS/AIAA Astrodynamics Specialist Conference*, AAS Paper 20-594, Springfield, VA, 2020.

[14] Friend, R. B., “Orbital Express Program Summary and Mission Overview,” *Proceedings of SPIE*, Vol. 6958, April 2008, Paper 695803. https://doi.org/10.1117/12.783792

[15] D’Amico, S., Ardaens, J.-S., Gaias, G., Benninghoff, H., Schlepp, 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, No. 6, 2013, pp. 1576–1595. https://doi.org/10.2514/1.59236

[16] 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, No. 2, 2014, pp. 439–451. https://doi.org/10.2514/1.61494

[17] 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, 2017, pp. 1337–1354. https://doi.org/10.2514/1.G003239

[18] 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

[19] Sanchez, H., McIntosh, D., Cannon, H., Pires, C., Sullivan, J., D’Amico, S., and O’Connor, B., “Starling-1: Swarm Technology Demonstration,” *Proceedings of the 32nd Small Satellite Conference*, Paper SSC18-VII-08, AIAA/USU, 2018.

[20] Miller, S., Adams, C., Alem, N., Cannon, H., Grashuis, R., Hendriks, T., Hwang, S., Iatauro, M., Pires, C., Kruger, J., et al., “Starling CubeSat Swarm Technology Demonstration Flight Results,” *Proceedings of the 38th Small Satellite Conference*, Paper SSC24-I-06, AIAA/USU, 2024.

[21] Sullivan, J., Koenig, A. W., Kruger, J., and D’Amico, S., “Generalized Angles-Only Navigation Architecture for Autonomous Distributed Space Systems,” *Journal of Guidance, Control, and Dynamics*, Vol. 44, No. 6, 2021, pp. 1087–1105. https://doi.org/10.2514/1.G005439

[22] Kruger, J., Koenig, A. W., and D’Amico, S., “The Starling Formation-Flying Optical Experiment (StarFOX): System Design, Experiment Design and Pre-Flight Verification,” *Journal of Spacecraft and Rockets*, Vol. 60, No. 6, 2023, pp. 1–23. https://doi.org/10.2514/1.A35598

[23] Kruger, J., Hwang, S., and D’Amico, S., “Starling Formation-Flying Optical Experiment: Initial Operations and Flight Results,” *Proceedings of the 38th Small Satellite Conference*, Paper SSC24-XI-01, AIAA/USU, 2024.

[24] Kruger, J., and D’Amico, S., “On-Orbit Performance and Lessons Learned for Autonomous Angles-Only Navigation of a Satellite Swarm,” *IEEE Aerospace Conference*, Inst. of Electrical and Electronics Engineers, New York, 2025. https://doi.org/10.1109/AERO63441.2025.11068739

[25] Vallado, D. A., and McClain, W. D., *Fundamentals of Astrodynamics and Applications*, 4th ed., Microcosm Press, Hawthorne, CA, 2013, pp. 145–158.

[26] D’Amico, S., “Autonomous Formation Flying in Low Earth Orbit,” Ph.D. Thesis, Delft Univ., Delft, The Netherlands, 2010, pp. 21–40.

[27] 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. https://doi.org/10.2514/1.G002409

[28] 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. https://doi.org/10.2514/1.G002822

[29] Sullivan, J., Lovell, T. A., and D’Amico, S., “Angles-Only Navigation for Autonomous On-Orbit Space Situational Awareness Applications,” *AAS/AIAA Astrodynamics Specialist Conference*, AAS Paper 18-468, Springfield, VA, 2018.

[30] Kruger, J., Wallace, K., Koenig, A. W., and D’Amico, S., “Autonomous Angles-Only Navigation for Spacecraft Swarms Around Planetary Bodies,” *IEEE Aerospace Conference*, Inst. of Electrical and Electronics Engineers, New York, 2021.

[31] Iiyama, K., Kruger, J., and D’Amico, S., “Autonomous Distributed Angles-Only Navigation and Timekeeping in Lunar Orbit,” *ION International Technical Meeting*, Inst. of Navigation, Manassas, VA, 2022, pp. 453–470. https://doi.org/10.33012/2022.18207

[32] Alfriend, K. T., edited by, *Spacecraft Formation Flying: Dynamics, Control, and Navigation*, Butterworth-Heinemann, New York, 2010, pp. 28–36. https://doi.org/10.1016/C2009-0-17485-8

[33] Montenbruck, O., and Gill, E., *Satellite Orbits: Models, Methods and Applications*, Springer, Berlin, 2012, pp. 53–55.
https://doi.org/10.1115/1.1451162

[34] Brouwer, D., "Solution of the Problem of Artificial Satellite Theory Without Drag," *Astronomical Journal*, Vol. 64, No. 1274, 1959, pp. 378–397.
https://doi.org/10.1086/107958

[35] 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, Springfield, VA, 2016.

[36] 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, 2017, pp. 2183–2200.
https://doi.org/10.2514/1.G002719

[37] Kruger, J., and D'Amico, S., "Autonomous Angles-Only Multitarget Tracking for Spacecraft Swarms," *Acta Astronautica*, Vol. 189, Dec. 2021, pp. 514–529.
https://doi.org/10.1016/j.actaastro.2021.08.049

[38] Vo, B.-N., Mallick, M., Bar-Shalom, Y., Coraluppi, S., Osborne, R., Mahler, R., and Vo, B. T., "Multitarget Tracking," *Wiley Encyclopedia of Electrical and Electronics Engineering*, edited by J. G. Webster, Springer International Publishing, Berlin, 2015, pp. 1–15.
https://doi.org/10.1002/047134608X.W8275

[39] Akondi, V., Roopashree, M. B., and Budihala, R. P., "Improved Iteratively Weighted Centroiding for Accurate Spot Detection in Laser Guide Star Based Shack Hartmann Sensor," *Atmospheric and Oceanic Propagation of Electromagnetic Waves IV*, edited by O. Korotkova, Vol. 7588, International Soc. for Optics and Photonics, SPIE, Bellingham, Washington, D.C., 2010, Paper 758806.
https://doi.org/10.1117/12.841331

[40] Wan, X., Wang, G., Wei, X., Li, J., and Zhang, G., "Star Centroiding Based on Fast Gaussian Fitting for Star Sensors," *Sensors*, Vol. 18, No. 9, 2018, Paper 2836.
https://doi.org/10.3390/s18092836

[41] 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

[42] Wertz, J. R., *Spacecraft Attitude Determination and Control*, Springer Science & Business Media, Berlin, 2012, pp. 410–435.
https://doi.org/10.1007/978-94-009-9907-7

[43] Blackman, S., "Multiple Hypothesis Tracking for Multiple Target Tracking," *IEEE Aerospace and Electronic Systems Magazine*, Vol. 19, No. 1, 2004, pp. 5–18.
https://doi.org/10.1109/MAES.2004.1263228

[44] Ester, M., Kriegel, H.-P., Sander, J., and Xu, X., "A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise," *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*, American Assoc. for Artificial Intelligence Press, Cambridge, MA, 1996, pp. 226–231.
https://doi.org/10.5555/3001460.3001507

[45] Thrun, S., Burgard, W., and Fox, D., *Probabilistic Robotics*, MIT Press, Cambridge, MA, 2005, p. 72.

[46] Ardaens, J.-S., and Gaias, G., "A Numerical Approach to the Problem of Angles-Only Initial Relative Orbit Determination in Low Earth Orbit," *Advances in Space Research*, Vol. 63, No. 12, 2019, pp. 3884–3899.
https://doi.org/10.1016/j.asr.2019.03.001

[47] Kulik, J., and Savransky, D., "State Transition Tensors for Passive Angles-Only Relative Orbit Determination," *33rd AAS/AIAA Space Flight Mechanics Meeting*, AAS Paper 23-161, Springfield, VA, 2023.

[48] Willis, M., and D'Amico, S., "Fast Angles-Only Relative Navigation Using Polynomial Dynamics," *Advances in Space Research*, Vol. 73, No. 11, 2024, pp. 5484–5500.
https://doi.org/10.1016/j.asr.2023.07.043

[49] Stacey, N., and D'Amico, S., "Autonomous Swarming for Simultaneous Navigation and Asteroid Characterization," *AAS/AIAA Astrodynamics Specialist Conference*, AAS Paper 18-448, Springfield, VA, 2018.

[50] Montenbruck, O., Kirschner, M., D'Amico, S., and Bettadpur, S., "E/I-Vector Separation for Safe Switching of the GRACE Formation," *Aerospace Science and Technology*, Vol. 10, No. 7, 2006, pp. 628–635.
https://doi.org/10.1016/j.ast.2006.04.001

[51] Hu, Y., Sharf, I., and Chen, L., "Distributed Orbit Determination and Observability Analysis for Satellite Constellations with Angles-Only Measurements," *Automatica*, Vol. 129, July 2021, Paper 109626.
https://doi.org/10.1016/j.automatica.2021.109626

C. D'Souza
Associate Editor