

The Starling Formation-Flying Optical Experiment (StarFOX): System Design and Pre-flight Verification

Justin Kruger *

Stanford University, Stanford, California, 94305

Adam W. Koenig[†]

Lockheed Martin, Sunnyvale, California, 94089

Simone D'Amico[‡]

Stanford University, Stanford, California, 94305

The Starling Formation-flying Optical eXperiment (StarFOX) is intended as the first on-orbit demonstration of autonomous distributed angles-only navigation for spacecraft swarms. StarFOX applies the angles-only Absolute and Relative Trajectory System (ARTMS), a navigation architecture consisting of three innovative algorithms: Image Processing, which identifies and tracks multiple targets in images from a single camera without a-priori relative orbit knowledge; Batch Orbit Determination, which autonomously initializes orbit estimates for visible swarm members; and Sequential Orbit Determination, which continuously refines the swarm state by fusing measurements from multiple observers exchanged over an inter-satellite link. Nonlinear dynamics and measurement models provide sufficient observability to estimate absolute orbits, relative orbits and auxiliary states using only bearing angles without maneuvers. StarFOX will be conducted using a four-CubeSat swarm as part of the NASA Starling mission, and simulations of experiment scenarios demonstrate that ARTMS meets mission performance requirements. Results indicate that mean bearing angle errors are below $35''$ (1σ), initial target range errors are below 20% of true separation, and steady-state range errors are below 2% (1σ). Absolute orbit estimation accuracy is on the order of 100 m. Hardware-in-the-loop tests display robust navigation under a variety of conditions, enabling autonomous, ubiquitous navigation with minimal ground interaction for future distributed missions.

Nomenclature

a = orbit semimajor axis

*Ph. D. Candidate, Department of Aeronautics and Astronautics, 496 Lomita Mall. Student Member AIAA.

[†]Senior GNC Engineer, Lockheed Martin, Sunnyvale, 1111 Lockheed Martin Way. Paper contributions were made as a Postdoctoral Scholar at Stanford University.

[‡]Associate Professor, Department of Aeronautics and Astronautics, 496 Lomita Mall. Associate Fellow AIAA.

α = azimuth bearing angle
 $\boldsymbol{\alpha}$ = orbit state vector
 δa = relative semimajor axis
 $\delta \boldsymbol{\alpha}$ = relative orbit state vector
 δe = magnitude of relative eccentricity vector
 δi = magnitude of relative inclination vector
 $\delta \lambda$ = relative mean longitude
 $\delta \boldsymbol{r}$ = relative position vector
 Δv = magnitude of maneuver delta-v
 e = orbit eccentricity
 ϵ = elevation bearing angle
 f = orbit true anomaly
 i = orbit inclination
 M = orbit mean anomaly
 Ω = orbit right ascension of the ascending node
 ω = orbit argument of periapsis
 \boldsymbol{P} = state covariance matrix
 ϕ = phase of relative inclination vector
 r = orbit radius
 \boldsymbol{r} = absolute position vector
 \boldsymbol{R} = measurement covariance matrix
 $\vec{\boldsymbol{R}}$ = rotation matrix
 \boldsymbol{S} = bearing angle measurement sensitivity matrix
 σ = standard deviation
 t = time
 \boldsymbol{t} = time vector
 θ = phase of relative eccentricity vector
 u = orbit mean argument of latitude
 \boldsymbol{x} = estimated state vector
 \boldsymbol{y} = bearing angle measurement vector
 z = bearing angle measurement batch

Subscripts

est = related to an estimated quantity
meas = related to a measured quantity
model = related to a modeled quantity
prior = related to an a-priori quantity
 o = observer spacecraft
 t = target object

Superscripts

\mathcal{I} = inertial reference frame
 \mathcal{R} = observer radial/tangential/normal reference frame
 \mathcal{V} = observer camera-aligned reference frame
 \mathcal{W} = observer velocity-aligned reference frame

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], as well as the proposed application of DSS to areas such as space domain awareness (SDA) [6] and on-orbit servicing, assembly and manufacturing (OSAM) [7]. However, robust navigation for DSS remains a technological challenge. The majority of DSS have been deployed in Earth orbit and their navigation systems therefore assume availability of external metrologies such as Global Navigation Satellite System (GNSS) signals and frequent contact with the ground. Systems outside of Earth orbit may instead 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 usage of 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 already 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 also equipped with an inter-satellite link (ISL), as can be implemented with typical radio frequency hardware, measurements from multiple observers can be shared and 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 may obtain

measurements to and navigate with respect to non-cooperative targets. However, bearing angles do not provide explicit target range information, which results in challenging observability conditions [9, 10]. Target range is only weakly observable and it is often difficult to simultaneously estimate absolute and relative orbit states [11].

These aspects have motivated extensive prior research into angles-only navigation for spacecraft, which may be subdivided into three primary tasks. First is the usage of sequential filtering techniques to perform angles-only state estimation. Earlier studies have focused on single-observer single-target scenarios, in which angles-only observability is particularly challenging. Nevertheless, it has been demonstrated that nonlinear filtering provides relative orbit observability without requiring maneuvers [10], and absolute orbit observability can potentially be achieved if observer and target orbits are sufficiently different [12, 13]. Recent work has explored usage of multiple cooperative observers to achieve complete DSS orbit observability [14, 15]. The second task is the usage of batch measurement techniques to perform initial angles-only orbit determination. Iterative numeric techniques have been proposed which display promising accuracy for estimation of relative orbits and target range with varying computational costs [11, 16]. Other authors have suggested computing relative orbit solutions using fewer measurements, via application of second-order polynomial models [17] or second-order state transition tensors [18] to resolve target range. The third task is target tracking and measurement assignment, or how bearing angles are obtained from images and consistently assigned to targets over time so that orbit determination can be performed. Common target tracking methods include global nearest neighbor (GNN), joint probabilistic data association (JPDA), multi-hypothesis tracking (MHT), random finite set (RFS), and machine learning (ML) approaches [19]. For in-orbit tracking specifically, it has been suggested to apply probability hypothesis density filtering [20] or domain-specific MHT [21].

While each of these algorithms shows varying degrees of promise, it is crucial to note that there is substantial difficulty in bridging the gap between individual algorithms that have been verified in simulation and a complete, integrated architecture suitable for flight which robustly performs all three tasks. Practical constraints including stringent mission requirements, software and hardware constraints, and system integration and validation must be overcome. With this in mind, two prior flight experiments have successfully demonstrated angles-only navigation in orbit, as documented in literature. In 2012, the Advanced Rendezvous using GPS and Optical Navigation (ARGON) experiment enabled the rendezvous of two SmallSats in low Earth orbit (LEO) from inter-satellite separations of 30 km to 3 km [22]. In 2016, the Autonomous Vision Approach Navigation and Target Identification (AVANTI) experiment similarly conducted a rendezvous between a mothership SmallSat and deployed picosatellite from separations of 13 km to 50 m [23]. However, these demonstrations were characterized by four major deficiencies that must be removed to meet the needs of future missions: 1) inability to accommodate multiple observers and multiple targets in the system, 2) lack of autonomy and reliance on accurate a-priori relative orbit information from the ground to initialize navigation, 3) reliance on external knowledge of the observer's absolute orbit (e.g. from a GNSS receiver) to maintain state convergence, and 4) reliance on frequent translational maneuvers to resolve the weakly observable inter-spacecraft separation.

To overcome these limitations and meet the navigation requirements of future DSS missions, this paper presents a high-level overview of the angles-only Absolute and Relative Trajectory Measurement System (ARTMS). ARTMS is a self-contained architecture first proposed in [14] that provides distributed, autonomous, scalable navigation capabilities for DSS orbiting an arbitrary central body. ARTMS consists of three core modules based on algorithms recently developed at Stanford’s Space Rendezvous Laboratory (SLAB): image processing (IMP) [21], batch orbit determination (BOD) [11], and sequential orbit determination (SOD) [14]. Each module is designed to operate with minimal a-priori information and exploit absolute and relative state knowledge as it becomes available. Overall, ARTMS provides real-time orbit estimates for the host spacecraft and each target detected by the onboard sensor as long as each observer is provided with an estimate of its orbit at a single epoch. Optional auxiliary states such as clock biases, ballistic coefficients and sensor biases can also be estimated using angles-only measurements. The only hardware requirements posed on a spacecraft with ARTMS are that it must have a VBS and an ISL. An optional absolute orbit metrology system such as a GNSS receiver can be included to improve navigation performance.

ARTMS will demonstrate these enhanced angles-only navigation capabilities during the Starling Formation-flying Optical eXperiment (StarFOX), as described and motivated in this paper. StarFOX is one of four experimental payloads of the NASA Starling mission in development at the NASA Ames Research Center [24]. Starling consists of four 6U CubeSats in LEO and aims to increase the readiness of four enabling technologies for spacecraft swarms: autonomous operations, communications and networking, maneuver planning and execution, and absolute and relative navigation. Starling is scheduled to launch in 2023 with a minimum mission duration of six months. A series of multi-day StarFOX experiment blocks is designed to explore ARTMS’ flexibility for angles-only navigation in both single-observer and distributed multi-observer scenarios, using either ground-assisted or autonomous state initializations, with or without GNSS availability, in three different swarm formations. The experiment campaign subsequently builds upon ARGON and AVANTI for a more wide-ranging and ambitious demonstration of spaceborne angles-only navigation.

To ensure that ARTMS is able to achieve StarFOX goals, a pre-flight testing and verification campaign is conducted using software- and hardware-in-the-loop (HIL) simulations. First, unification of ARTMS algorithms into the complete flight software is motivated by extensive Monte-Carlo characterization of individual IMP, BOD and SOD module performance in the context of StarFOX mission requirements. Next, performance of the complete ARTMS architecture is assessed using high-fidelity simulations representative of planned StarFOX experiments, for varying swarm geometries and measurement conditions. HIL elements include usage of a Nano Star Tracker (NST) from Blue Canyon Technologies [8] as stimulated by SLAB’s Optical Stimulator (OS) testbed [25] to generate image measurements with realistic noise characteristics; usage of a Tyvak Endeavour CubeSat processor to assess ARTMS computation costs; and measurement error characterization for each NST flight unit when integrated with the spacecraft.

After this introduction, key modeling assumptions are provided in Section II. A high-level overview of ARTMS and its constituent modules is provided in Section III. Section IV introduces and motivates the StarFOX flight experiment.

Section V presents Monte-Carlo performance characterization of individual ARTMS modules, followed by validation of the integrated ARTMS flight software. Section VI presents HIL test results. Conclusions are given in Section VII.

II. Modeling Preliminaries

A. Measurement Model

ARTMS produces angles-only measurements by computing the time-tagged bearing angles to objects detected in VBS images. Two rotating coordinate frames are defined. First, consider the radial/tangential/normal (RTN) frame of an observer spacecraft, denoted \mathcal{R} . It is centered on and rotates with the observer and consists of orthogonal basis vectors $\hat{\mathbf{x}}^{\mathcal{R}}$ (directed along the observer's absolute position vector); $\hat{\mathbf{z}}^{\mathcal{R}}$ (directed along the observer's orbital angular momentum vector); and $\hat{\mathbf{y}}^{\mathcal{R}} = \hat{\mathbf{z}}^{\mathcal{R}} \times \hat{\mathbf{x}}^{\mathcal{R}}$ [26]. Similarly, define a frame \mathcal{W} using $\hat{\mathbf{y}}^{\mathcal{W}}$ (directed along the observer's velocity vector); $\hat{\mathbf{z}}^{\mathcal{W}} = \hat{\mathbf{z}}^{\mathcal{R}}$; and $\hat{\mathbf{x}}^{\mathcal{W}} = \hat{\mathbf{y}}^{\mathcal{W}} \times \hat{\mathbf{z}}^{\mathcal{W}}$. \mathcal{W} only differs from \mathcal{R} by a rotation of the observer flight path angle ϕ_f about $\hat{\mathbf{z}}^{\mathcal{R}}$ with $\phi_f \approx 0$ in near-circular orbits [26]. Finally, define the observer VBS coordinate frame \mathcal{V} consisting of orthogonal basis vectors $\hat{\mathbf{x}}^{\mathcal{V}}, \hat{\mathbf{y}}^{\mathcal{V}}, \hat{\mathbf{z}}^{\mathcal{V}}$ where $\hat{\mathbf{z}}^{\mathcal{V}} = \hat{\mathbf{x}}^{\mathcal{V}} \times \hat{\mathbf{y}}^{\mathcal{V}}$ is aligned with the camera boresight. The VBS may be pointed as necessary to keep targets in the field of view (FOV). In the case of StarFOX, it is chosen to always point the camera boresight in the velocity or anti-velocity direction for simplicity, such that $\hat{\mathbf{z}}^{\mathcal{V}}$ is aligned with $\pm \hat{\mathbf{y}}^{\mathcal{W}}$. Figure 1 illustrates this scenario.




Fig. 1 Definition of coordinate frames and bearing angle geometry with VBS pointing in the $-\hat{\mathbf{y}}^{\mathcal{W}}$ direction.

Bearing angles consist of azimuth and elevation $[\alpha, \epsilon]^T$ and subtend the line-of-sight vector $\delta \mathbf{r}^{\mathcal{V}}$ from the observer to the target. The measurement model \mathbf{y} for the bearing angles from observer to target is described by [14]

$$\delta \mathbf{r}^{\mathcal{V}} = \mathbf{r}_t^{\mathcal{V}} - \mathbf{r}_o^{\mathcal{V}} = [\delta r_x^{\mathcal{V}}, \delta r_y^{\mathcal{V}}, \delta r_z^{\mathcal{V}}]^T \quad (1)$$

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

System measurements and states are also referenced with respect to an inertial reference frame centered on an arbitrary central body, denoted \mathcal{I} (commonly a planet-centered, moon-centered or Sun-centered frame). The inertial frame applied during StarFOX is the Earth-centered inertial J2000 frame. Bearing angles are related to the inertial frame by rotating $\delta\mathbf{r}^{\mathcal{V}}$ into \mathcal{I} , as per $\delta\mathbf{r}^{\mathcal{I}} = {}^{\mathcal{V}}\overrightarrow{\mathbf{R}}^{\mathcal{I}} \delta\mathbf{r}^{\mathcal{V}}$ where ${}^{\mathcal{V}}\overrightarrow{\mathbf{R}}^{\mathcal{I}}$ 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 [22]. Other relevant rotations ${}^{\mathcal{R}}\overrightarrow{\mathbf{R}}^{\mathcal{I}}$ and ${}^{\mathcal{W}}\overrightarrow{\mathbf{R}}^{\mathcal{I}}$ can be computed using the observer's absolute orbit estimate. Note that in practice, VBS measurement availability is affected by optical visibility constraints including target eclipse periods, sensor sun-blinding, sensor FOV, and target visual magnitude versus sensor sensitivity.

B. System State

Because angles-only navigation is known to be characterized by weak observability [9], proper selection of the state parameterization is particularly important to maximize accuracy of state estimates and robustness to errors in a-priori information. Specifically, the state definition should be selected to meet two objectives: 1) separate weakly and strongly observable terms to enable use of separate estimation techniques, and 2) maximize accuracy and computational efficiency of orbit propagation. To meet both of these objectives, the system state is defined using quasi-nonsingular absolute orbit elements (OE) and relative orbital elements (ROE). The absolute orbit α is defined as

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

where e_x and e_y are components of the eccentricity vector. The relative orbit of each target detected by the onboard sensor, denoted $\delta\alpha$, is described by the quasi-nonsingular ROE adopted by D'Amico [27]. Each of these ROE is function of the orbit elements of the target and observer as given by

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

where $(\delta e_x, \delta e_y)$ are components of the relative eccentricity vector and $(\delta i_x, \delta i_y)$ are components of the relative inclination vector. Section II.C describes the relationships between OE, ROE, and target relative motion in the $\hat{\mathbf{x}}^{\mathcal{R}}\text{-}\hat{\mathbf{y}}^{\mathcal{R}}$ (radial-tangential) and $\hat{\mathbf{x}}^{\mathcal{R}}\text{-}\hat{\mathbf{z}}^{\mathcal{R}}$ (radial-normal) planes.

These state definitions have been previously studied in literature, resulting in development of accurate analytical dynamics models (e.g. [28] and [29]). 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 [30]. More importantly, the weakly observable range to each target is primarily captured by the $\delta\lambda$ term in most relative motion geometries, especially for loose satellite formations and swarms. 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 [11]. Combined, these properties enable accurate and computationally efficient estimation algorithms with minimal reliance on a-priori information, as described in Section III.

ARTMS also possess the capacity to estimate auxiliary state components, such as differential clock offsets and drift rates between observers, differential ballistic coefficients, and sensor biases [14, 31–33], and may furthermore operate using fully nonsingular orbit elements in equatorial orbits [34]. However, only orbit estimation will be performed during the initial StarFOX experiment phase. This is motivated by a desire to reduce computation costs and system complexity. Furthermore, navigation errors due to mismodeled auxiliary states are expected to be minimal during StarFOX: clocks will be synchronized with GPS time, the proposed orbit will result in minimal atmospheric drag, and sensors have been calibrated before flight. For conciseness, auxiliary state estimation is not discussed here, but will potentially be explored as part of a StarFOX extended mission phase.

C. Dynamics Model

ARTMS propagates the absolute orbits of system objects using fourth-order Runge-Kutta integration of the GVE. For state α , the osculating OE of each spacecraft evolve according to

$$\dot{\alpha} = G(\alpha)\mathbf{d}^{\mathcal{R}} \quad (5)$$

$$\mathbf{d}^{\mathcal{R}} = \begin{bmatrix} d_x^{\mathcal{R}} & d_y^{\mathcal{R}} & d_z^{\mathcal{R}} \end{bmatrix}^{\top} \quad (6)$$

where $G \in \mathbb{R}^{6 \times 3}$ is the GVE matrix [35] and $\mathbf{d}^{\mathcal{R}}$ is the perturbing acceleration expressed in \mathcal{R} . Depending on the orbit regime, common perturbations include spherical harmonic gravity terms, atmospheric drag, third-body gravity and solar radiation pressure (SRP) [36]. Analytic dynamics models for the mean OE which include the effects of J_2 gravity are used within ARTMS when computational efficiency is paramount. Specific dynamics models and sources of perturbing accelerations applied within ARTMS are detailed in Section V.

A useful aspect of the ROE is that they provide geometric intuition regarding target relative motion. Before examining these dynamics, it is beneficial to define curvilinear coordinates which capture the effects of orbit curvature with improved accuracy. The curvilinear position vector of a target in the observer’s RTN frame is defined as $\delta\mathbf{r}_{\text{curv}}^{\mathcal{R}} = (\delta r, a_o\Theta, a_o\Phi)$

[37]. Here, δr , Θ , Φ are target-observer differences in orbit radii, angular in-plane separations and angular out-of-plane separations respectively. The curvilinear representation can be mapped back to rectilinear coordinates via

$$\delta \mathbf{r}_{\text{rect}}^{\mathcal{R}} = \begin{bmatrix} (a_o + \delta r) \cos \Theta \cos \Phi - a_o \\ (a_o + \delta r) \sin \Theta \cos \Phi \\ (a_o + \delta r) \sin \Phi \end{bmatrix} \quad (7)$$

There exists a subsequent linear mapping between the ROE and $\delta \mathbf{r}_{\text{curv}}^{\mathcal{R}}$, first shown for near-circular orbits in [27]. The mapping was extended to eccentric orbits by defining the eccentric ROE $\delta \alpha^* = (\delta a, \delta \lambda^*, \delta e_x^*, \delta e_y^*, \delta i_x, \delta i_y)$, which revert to traditional, quasi-nonsingular ROE for $e_o \approx 0$ [38]. The resulting mapping is

$$\delta \mathbf{r}^{\mathcal{R}} \approx r_o \begin{bmatrix} \delta a - \frac{e_o}{2} \delta e_x^* - \delta e^* \left(\cos(f_o - \phi^*) + \frac{e_o}{2} \cos(2f_o - \phi^*) \right) \\ \delta \lambda^* + \delta e^* \left(2 \sin(f_o - \phi^*) + \frac{e_o}{2} \sin(2f_o - \phi^*) \right) \\ \delta i \sin(f_o + \omega_o - \theta) \end{bmatrix} \quad (8)$$

Figure 2 presents relative motion in RTN for this mapping for small target separations. Components of oscillatory motion produced by the 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. 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.




Fig. 2 Components of target relative motion in the $\hat{x}^{\mathcal{R}}-\hat{y}^{\mathcal{R}}$ (RT) and $\hat{x}^{\mathcal{R}}-\hat{z}^{\mathcal{R}}$ (RN) planes [14].

III. ARTMS Architecture

A. Overview

To simplify subsequent discussions, the following terminologies are adopted. The “observer” refers to the spacecraft hosting the instance of ARTMS being discussed. A “remote observer” is another spacecraft hosting ARTMS providing measurements over the 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” (or DSS) refers to the entire distributed space system, consisting of all involved observers and targets. Targets may be non-cooperative objects or spacecraft which do not actively assist navigation, or may be cooperative remote observers themselves. Figure 3 presents a notional illustration of a four-spacecraft DSS using ARTMS.




Fig. 3 Notional illustration of ARTMS observers and targets for a four-spacecraft system (not to scale).

A high-level overview of ARTMS as implemented in StarFOX flight software is shown in Figure 4. ARTMS

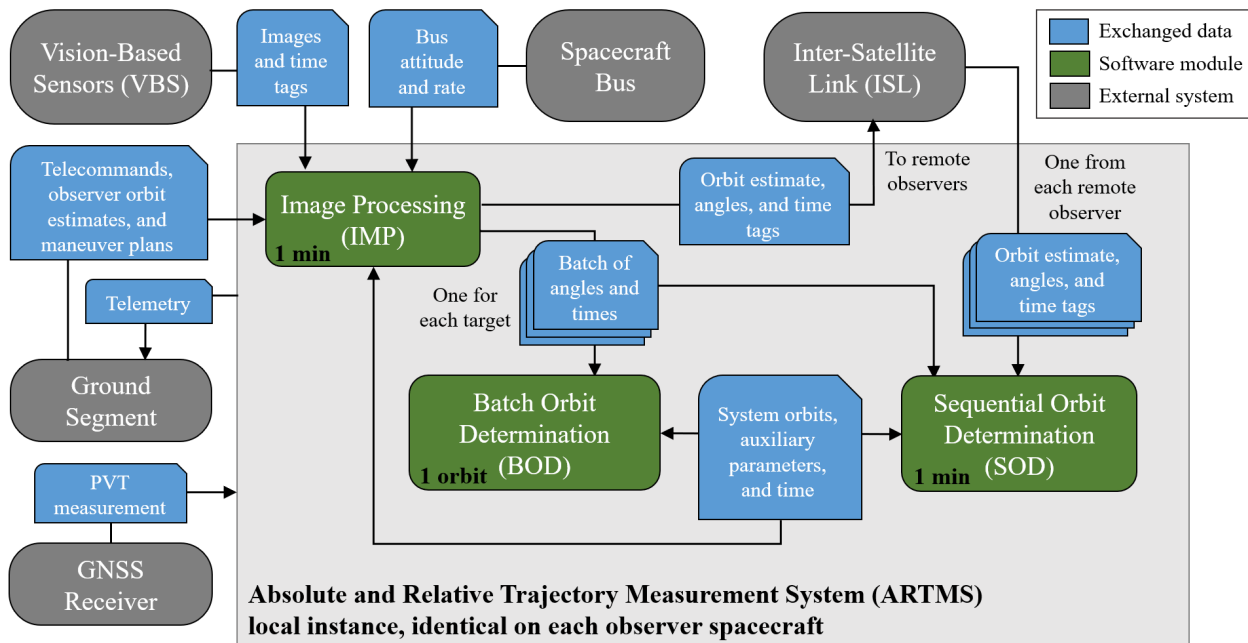


Fig. 4 A high-level overview of the ARTMS flight architecture.

is divided into three main modules: image processing (IMP), batch orbit determination (BOD) and sequential orbit determination (SOD). It interfaces with the onboard VBS, the ISL, the ground segment, and an onboard GNSS receiver (if available). The VBS provides time-tagged raw images, which are processed to obtain inertial bearing angles to target RSO. The ISL communicates orbit estimates and angles 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. This information is used by all ARTMS modules. Similarly, the ground segment receives telemetry from all modules. If the spacecraft is equipped with a GNSS receiver, it provides position/velocity/time (PVT) navigation solutions to all modules, to replace less timely orbit estimates from the ground. These absolute orbit estimates are necessary for each observer to initialize navigation.

The IMP module uses a coarse estimate of the observer's absolute orbit and VBS images to produce batches of bearing angle measurements and corresponding uncertainties to all visible targets without any 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 coarse estimate of the observer's absolute orbit and the batches of bearing angles provided by 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. ballistic coefficients or differential clock offsets). 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 reduce the computation cost of tracking detected targets.

ARTMS 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.

B. Assumptions and Limitations

The version of ARTMS to be flown during StarFOX removes many key limitations of prior angles-only flight experiments and is a significant advance over existing fully-integrated architectures. Nevertheless, ARTMS still requires three key assumptions to operate.

First, during StarFOX experiments featuring cooperative spacecraft, it is assumed that the ground provides a unique identification number (ID) and a-priori orbit estimate for each DSS member. Each observer can then utilize its a-priori orbit estimate to initialize navigation (but may not have access to the a-priori orbits estimates of other DSS members). Moreover, it is assumed maneuvers performed by the swarm are cooperative, such that the ID number, execution time,

burn time, magnitude, and direction of every swarm maneuver is provided to ARTMS. Future work will explore detection and classification of unknown target maneuvers via uncertainty-aware filtering and outlier detection techniques.

Second, StarFOX formations have been specifically designed to ensure targets remain consistently in the FOV of the observer. This may not be true for more general DSS geometries. If a-priori target state information is available, it is assumed observers have the ability to align their VBS to track target orbits of interest. If a-priori state information is unavailable, it is assumed targets cross the FOV with sufficient regularity to at minimum initialize tracking. As detailed in Section III.C, this implies a minimum of four target sightings in successive images. Future work will investigate the optimization of observer attitudes to ensure target measurements are obtained consistently, as well as the usage of wide-FOV cameras and variable measurement frequencies for detection of fast-moving targets across larger areas.

Third, the current structure of ARTMS possesses limited scalability to very large numbers of tracked RSO. Distribution and decentralization arise from hardware limitations, in that different observers may only be able to obtain measurements of specific subsets of targets due to limitations on VBS FOV, detectable visual magnitude, and radio range. However, each ARTMS observer automatically attempts to track all locally-detected RSO and broadcasts its measurements to all remote observers. Benefits of this approach include its simplicity and redundancy for flight experiments because it ensures the maximum number of measurements is available if one or more observers fails. Conversely, if there are N_t targets in the FOV where N_t is large, or N_o communicating observers where N_o is large, on-board navigation may become computationally intractable. This is especially important for the IMP module which applies multi-hypothesis tracking computations that scale quickly with N_t . Future versions of ARTMS will explore more intelligent methods of decentralized computation, such that observers can autonomously divide navigation and tracking tasks between active ARTMS instances in a more optimal fashion.

C. Image Processing

The objective of the IMP module is to produce batches of time-tagged bearing angle measurements to each target using a coarse estimate of the observer's orbit and images provided by the onboard VBS. This is accomplished in two phases. First, each incoming image is processed and reduced to a set of inertial bearing angles that may correspond to RSO. Second, these candidate bearing angles are used to track known targets and detect new targets using an approach inspired by multi-hypothesis tracking (MHT) [19]. Figure 5 presents an example image input with visible point sources. Typical sample rates for the VBS are 1-2 minutes in LEO.

The first phase of IMP uses a set of algorithms developed previously in literature. First, a Gaussian grid centroiding algorithm [39] is used to simplify the raw image into a list of pixel cluster centroids. Second, these centroids are converted to unit vectors in the VBS frame using a calibrated VBS model. Next, the pyramid star identification algorithm [40] is applied to remove stellar objects (SO) from the list of pointing vectors. The pointing vectors to identified stars in the inertial and VBS frames are used to compute the VBS attitude using the q-method [41]. The inertial unit vectors to

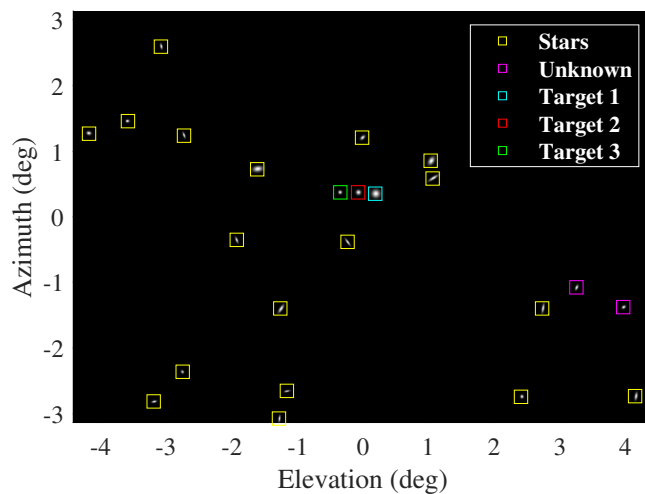


Fig. 5 Example input image with point sources detected and identified by IMP.

targets typically vary quickly in time due to orbital motion. Thus, objects with inertial unit vectors that remain similar between images are considered to be uncatalogued SO and removed. A similarity threshold of $20''$ allows for the effects of nominal bearing angle errors. Similarly, objects with unit vectors in \mathcal{V} that remain similar between images are considered to be camera hotspots and removed. A similarity threshold of $5''$ prevents mis-identifying targets with near-zero relative motion. The remaining minimalistic set of inertial unit vectors corresponds to known targets or other unknown objects in the FOV.

In the second phase, these measurements must be assigned to currently-tracked targets or used to initialize tracks for unidentified targets. To accomplish this, IMP employs the Spacecraft Angles-only MULTitarget tracking System (SAMUS) algorithm [21]. SAMUS is agnostic to orbit eccentricity and requires only 1) coarse absolute orbit knowledge of the observer and 2) knowledge of the magnitudes and execution times of swarm maneuvers (but not which maneuvers correspond to which targets). No knowledge of the number of targets or their relative orbits is needed. SAMUS is designed to meet the constraints of risk-averse angles-only navigation in space, i.e. to achieve close to 100% measurement assignment precision with low measurement frequencies and limited computational resources. It applies the core concept of MHT in that as measurements arrive, several simultaneous hypotheses are maintained as to how they can be associated into target tracks. The correct hypothesis is converged to over time as more information is received.

Apart from MHT, several other families of multitarget tracking algorithms were considered for SAMUS, including GNN, JPDA, RFS, and ML approaches. Each has advantages and disadvantages. Briefly, GNN is computationally simple but susceptible to poor performance when targets are not well-separated [42]. JPDA, though demonstrably accurate in many scenarios, generally assumes a known number of targets [43]. Both are then non-ideal in that targets may not be well-separated and the number of visible targets may be unknown. RFS techniques are more recent, with

many promising varieties under development; conversely, they are somewhat less proven and approximations are needed for real-time usage [19]. ML approaches for data association have also become increasingly popular but there are difficulties in generating training data representative of the space environment [44]. In contrast, MHT is a mature, theoretically optimal approach which performs well for low signal-to-noise ratios [42]. Its primary disadvantage is that the number of hypothesis may grow very quickly, resulting in a need to frequently and heuristically trim unlikely hypotheses to achieve real-time computation (particularly problematic for low-powered spacecraft processors).

In comparison to naive MHT, SAMUS achieves vastly improved efficiency by leveraging domain-specific knowledge to develop new hypothesis management criteria. Figure 6 presents an overview of core SAMUS operations, where dashed outlines denote steps that only occur at relevant epochs. Its output is a set of target tracks, which each consist of




Fig. 6 SAMUS algorithm summary and core sequence of operations.

bearing angles measured by the local observer, separated in time.

The first operation is to detect and initialize possible new target tracks by applying the Density-Based Spatial Clustering of Applications with Noise (DBSCAN) algorithm to unidentified points in the most recent four images [45]. During StarFOX and similar swarm scenarios, target velocities are low compared to other objects in the FOV, because targets are in similar absolute orbits to observers. Thus, target measurements form clusters in \mathcal{V} . A new SAMUS target is initialized if DBSCAN detects a cluster of at least four unidentified bearing angles within some small radius ϵ_D .

Subsequent operations apply Equation 8, which maps OE and ROE to target relative positions in \mathcal{R} . On the right hand side of Equation 8, true anomaly f_o is the only quickly-varying term, whereas other terms vary on longer timescales in the presence of perturbations such as J_2 in LEO [27]. The ‘measurement transform’ step in Figure 6 recalls that system members in formation-flying scenarios remain in close proximity in inertial space, such that each is affected similarly by perturbations. By synchronously differencing the bearing angles of different targets – in essence, using one track as a virtual, moving origin for another – perturbation effects are approximately cancelled between targets, recovering motion as per Equation 8. Target motion is then parametric in f_o with a known periodic form. Even if specific ROE are unknown, this model can be leveraged to assess bearing angle tracks.

The ‘measurement prediction’ operation rearranges the radial and normal components of Equation 8 so that a track

motion model may be fitted to its measurements, as per

$$\begin{bmatrix} \epsilon \\ \alpha \end{bmatrix}^{\mathcal{R}} \approx \frac{r_o}{a_o} \begin{bmatrix} x_1 - x_2(\cos(f_o - x_3) + \frac{e_o}{2} \cos(2f_o - x_3)) \\ x_4 + x_5 \sin(f_o + \omega - x_6) \end{bmatrix} \quad (9)$$

where $x_{1,\dots,6}$ are scaled ROE equivalents in bearing angle space [21]. Terms $r_o, a_o, f_o, e_o, \omega_o$ are computed from the observer orbit estimate. If a track consists of at least three bearing angles at different times, the six unknowns $x_{1,\dots,6}$ can be solved for via least squares. Measurements in future epochs can then be predicted by propagating the observer orbit estimate and applying the fitted model, and goodness of fit can be assessed via fitting residuals.

In the ‘track propagation’ step, SAMUS applies a set of kinematic rules (derived from the parametric motion model) to assess which tracks are physically reasonable. Figure 7 presents relevant track quantities in the two-dimensional




Fig. 7 Target track quantities in the two-dimensional bearing angle plane.

bearing angle plane. Vectors \vec{v}_k connect successive bearing angles, where the subscript k refers to the current tracking epoch. Each \vec{v}_k has magnitude d_k , phase ζ_k , and inter-vector angle ψ_k . Briefly, the rules are summarised as:

- 1) Target track velocities in the two-dimensional bearing angle plane must be below a user-specified maximum, as derived from orbit dynamics, the image measurement frequency, and constraints on the allowed magnitude of target relative motion, such that $d_k < d_{k,\max}$.
- 2) Target track velocities must be consistent over time such that the ratio of track velocities between successive images is close to one, as per $|v_k/v_{k-1}| \approx 1$.
- 3) Tracks should generally not feature acute angles, such that $\psi_k \approx \pi$.
- 4) Tracks should turn in a consistent direction such that they follow an approximately elliptical path that does not change significantly on the timescale of image-to-image tracking, 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 user-specified tolerance, as per $\|(\alpha, \epsilon)_k - (\alpha_p, \epsilon_p)_k\|_2 < d_{p,\max}$.

Tracks which do not pass all rules are considered unlikely and are removed, which increases the efficiency of MHT.

When multiple tracks pass all rules, SAMUS scores propagated tracks via a set of criteria which assess how well each fulfills the expectations of Equation 8 and their prior motion [21]. These criteria include:

- 1) The size of summed residuals from track fitting.
- 2) The difference between the newest predicted and assigned measurements, i.e. $\|(\alpha, \epsilon)_k - (\alpha_p, \epsilon_p)_k\|_2$.
- 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 ϕ_k to its predicted value and mean values, i.e. $|\phi_k - \phi_p|$ and $|\phi_k - \phi_{\text{mean}}|$.
- 5) The preference for tracks to possess lower velocities and more gradual turns, i.e. $|d_k|$ and $1/|\phi_k|$.

Lower scores are preferred. In contrast to traditional MHT methods which often rely on a single statistical metric for scoring, SAMUS aims to be more robust. Often, target 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, but consensus from a larger set of metrics is able to support the correct choice over time. Additionally, scoring does not require probabilistic estimates of false alarm densities or target decay rates which are not easily obtainable for spacecraft.

Surviving hypotheses are then managed and pruned according to traditional MHT methods [46]. If maneuvers are expected during tracking, they are matched to targets by examining the qualitative change in their fitted x_1, \dots, x_6 values pre- and post-maneuver, and comparing this to expected qualitative changes in ROE from a state transition matrix. Finally, tracks must also be re-initialized in the event of on-orbit measurement interruptions such as eclipse periods. To connect shorter tracks on either side of an interruption, the linear system fit of Equation 9 is computed for every track pair. The combination of compatible pairs which produces the least fitting residuals is chosen as output.

In scenarios where a-priori relative state estimates are available from SOD, the process is simplified. Track measurement predictions and an associated uncertainty region are provided via an unscented transform from the ROE state to bearing angles. The uncertainty region is applied for track gating and the Mahalanobis distance between predicted and assigned measurements is employed for track scoring. Note that the linearizations present in Equation 8 imply SAMUS is currently most appropriate for inter-spacecraft separations up to several hundred kilometers – as encountered during StarFOX and swarm scenarios – and not necessarily constellation at very large separations. Improved motion models such as those in [17] will be leveraged to explore the constellation case.

D. Batch Orbit Determination

The BOD module must produce orbit estimates for visible targets with sufficient accuracy to initialize the SOD module using only a coarse estimate of the observer’s orbit and batches of bearing angles to each target from IMP. State estimation is accomplished using a spaceborne algorithm [11] inspired by the work of Ardaens & Gaias [16] that sequentially estimates the six ROE of each target while simultaneously refining the estimate of the observer’s semimajor axis, using batches of bearing angle measurements. Typical measurement collection periods are 1-2 orbits for a total batch of $N_y \in [50, 200]$ measurements per target. The batch estimation problem is somewhat ill-conditioned due to

the weak angles-only observability of $\delta\lambda$. As a result, $\delta\lambda$ is estimated using a line sampling procedure whereas the other strongly-observable ROE are estimated using iterative batch least squares. Note that currently, BOD only utilizes measurements from the onboard camera; although more geometric information and improved $\delta\lambda$ observability would be gained by extending the method to multiple observers, additional challenges are faced in regards to algorithmic complexity, communications bandwidth, and computational requirements. Multi-observer BOD capabilities are therefore proposed for a StarFOX extended mission, and the following algorithm focuses on single-observer BOD.

Consider a model $\mathbf{y}^I(\mathbf{x}(t_{\text{est}}), t)$ that provides bearing angles at time t as a function of system state $\mathbf{x}(t_{\text{est}})$ at estimation epoch t_{est} . It is useful to partition \mathbf{x} into components estimated and components provided a-priori, with $\mathbf{x} = [\mathbf{x}_{\text{est}}^T, \mathbf{x}_{\text{prior}}^T]^T$. For N_x estimated state components, $\mathbf{x}_{\text{est}} \in \mathbb{R}^{N_x}$, where StarFOX applies $\mathbf{x}_{\text{est}} = [\delta a, \delta\lambda, \delta e_x, \delta e_y, \delta i_x, \delta i_y, a_o] \in \mathbb{R}^7$.

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

$$\mathbf{z}(\mathbf{x}(t_{\text{est}}), \mathbf{t}) = \begin{pmatrix} \mathbf{y}^I(\mathbf{x}(t_{\text{est}}), t_1) \\ \vdots \\ \mathbf{y}^I(\mathbf{x}(t_{\text{est}}), t_{N_y}) \end{pmatrix} \quad (10)$$

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

$$\mathbf{S}_{\text{est}}(\mathbf{x}(t_{\text{est}}), \mathbf{t}) = \left. \frac{\partial \mathbf{z}(\mathbf{x}(t_{\text{est}}), \mathbf{t})}{\partial \mathbf{x}_{\text{est}}} \right|_{\mathbf{x}(t_{\text{est}})} \quad \mathbf{S}_{\text{prior}}(\mathbf{x}(t_{\text{est}}), \mathbf{t}) = \left. \frac{\partial \mathbf{z}(\mathbf{x}(t_{\text{est}}), \mathbf{t})}{\partial \mathbf{x}_{\text{prior}}} \right|_{\mathbf{x}(t_{\text{est}})} \quad (11)$$

with $\mathbf{S}_{\text{est}} \in \mathbb{R}^{2N_y \times N_x}$.

For each target, the estimation process applies the following procedure. First, the user selects a set of samples of $\delta\lambda$. The $\delta\lambda$ state space may be bounded by the expected minimum target separation and the maximum target range physically observable by the VBS. The sampling interval may be selected with reference to on-board computational resources or the desired $\delta\lambda$ output accuracy. StarFOX divides the expected state space for $\delta\lambda$ into 100-300 intervals in the positive and negative directions with sampling intervals of 1-2 km.

For each $\delta\lambda$ sample, \mathbf{x}_{est} is initialized as follows. For the first sample, all ROE except $\delta\lambda$ are set to zero, and a is set to the value from the observer's absolute orbit estimate. For subsequent samples, ROE are initialized as

$$\delta\alpha^+ = \delta\alpha^- \frac{\delta\lambda^+}{\delta\lambda^-} \quad a^+ = a^- \quad (12)$$

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

At each iteration, a modeled batch of bearing angle measurements $\mathbf{z}_{\text{model}} \in \mathbb{R}^{2N_y}$ is computed by propagating the estimated state vector \mathbf{x}_{est} to each measurement epoch. The difference between measured and modeled angles is then

$$\Delta \mathbf{z} = \mathbf{z}_{\text{meas}} - \mathbf{z}_{\text{model}}(\mathbf{x}(t_{\text{est}}), \mathbf{t}) \quad (13)$$

where $\mathbf{z}_{\text{meas}} \in \mathbb{R}^{2N_y}$ is the batch of true measurements from IMP. Subsequently, the state estimate update

$$\Delta \mathbf{z} = \mathbf{S}_{\text{est}}(\mathbf{x}(t_{\text{est}}), \mathbf{t}) \Delta \mathbf{x}_{\text{est}} \quad (14)$$

$$\mathbf{x}_{\text{est}} \leftarrow \mathbf{x}_{\text{est}} + \Delta \mathbf{x}_{\text{est}} \quad (15)$$

can be performed, where $\mathbf{S}_{\text{est}}(\mathbf{x}(t_{\text{est}}), \mathbf{t})$ is computed via linearization of the bearing angle measurement model and the unknown $\Delta \mathbf{x}_{\text{est}}$ is solved for via least squares. State refinement continues until a user-specified iteration limit is reached or $\Delta \mathbf{x}_{\text{est}}$ is smaller than a user-specified convergence threshold.

After refinement has been completed for each $\delta \lambda$ sample, the final state estimate \mathbf{x}_f is selected as the converged sample which produced the smallest measurement residual vector $\Delta \mathbf{z}$. This leads to an optimization objective of

$$\min_{\mathbf{x}_{\text{est}}} \Delta \mathbf{z}^\top \Delta \mathbf{z} \quad (16)$$

for which $\mathbf{x}_{\text{est}} = \mathbf{x}_f$ is the optimal solution. Figure 8 is a notional illustration of 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 [16].




Fig. 8 Behavior of converged measurement residuals for different $\delta \lambda$ samples in the BOD module.

It is then necessary to estimate the associated initial state uncertainty. The equation applied is [11]

$$\mathbf{R}_{\text{est}} = \mathbf{S}_{\text{est}} \mathbf{P}_{\text{est}} \mathbf{S}_{\text{est}}^\top$$

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

$$\begin{aligned} \mathbf{R}_{\text{est}} &= \mathbf{R}_{\text{sens}} + \mathbf{R}_{\text{prior}} + \mathbf{R}_{\text{dyn}} \\ &= \begin{bmatrix} \mathbf{R}_{\text{meas}} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbf{R}_{\text{meas}} & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \mathbf{R}_{\text{meas}} \end{bmatrix} + \mathbf{S}_{\text{prior}}(\mathbf{x}_f) \mathbf{P}_{\text{prior}} \mathbf{S}_{\text{prior}}(\mathbf{x}_f) + \begin{bmatrix} \mathbf{R}_{\text{proc}} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbf{R}_{\text{proc}} & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \mathbf{R}_{\text{proc}} \end{bmatrix} \end{aligned}$$

The \mathbf{R}_{sens} term encompasses uncertainties due to sensor noise; the $\mathbf{R}_{\text{prior}}$ term encompasses uncertainties stemming from a-priori information; and the \mathbf{R}_{dyn} term encompasses uncertainties due to systematic discrepancies between the on-board and ground truth dynamics models. Each term can be computed. The sensor noise associated with individual measurements, $\mathbf{R}_{\text{meas}} \in \mathbb{R}^{2 \times 2}$, is estimated via

$$\mathbf{R}_{\text{meas}} = \frac{1}{N_y} \sum_{j=1}^N (\mathbf{z}_{\text{meas}}(t_j) - \mathbf{y}^T(\mathbf{x}_f, t_{\text{est}}, t_j)) (\mathbf{z}_{\text{meas}}(t_j) - \mathbf{y}^T(\mathbf{x}_f, t_{\text{est}}, t_j))^\top \quad (17)$$

This estimation of sensor noise using the post-fit measurement residuals allows operation even when an accurate a-priori model of sensor noise is unavailable. The uncertainty of a-priori state information, $\mathbf{P}_{\text{prior}}$, is assumed to be provided. The measurement uncertainty due to process noise, $\mathbf{R}_{\text{proc}} \in \mathbb{R}^{2 \times 2}$, can be estimated by propagating the time-varying orbit covariance due to unmodeled RTN accelerations to each measurement epoch, then applying an unscented transform to convert the orbit covariance to a bearing angle covariance, following the process in [33]. In this fashion, time-varying uncertainties due to propagation times from each measurement epoch to the BOD estimation epoch are accounted for. However, instead of computing \mathbf{R}_{dyn} , StarFOX applies the empirical approximation $\mathbf{R}_{\text{est}} = N_y \mathbf{R}_{\text{sens}} + \mathbf{R}_{\text{prior}}$ [11]. It has been found in practice that this factor of N_y leads to sufficient consistency between BOD estimation errors and uncertainties without requiring additional computation. The initial state uncertainty can then be estimated from known terms via $\mathbf{P}_{\text{est}} = \mathbf{S}_{\text{est}}^* \mathbf{R}_{\text{est}} \mathbf{S}_{\text{est}}^{*\top}$, where $*$ denotes a pseudoinverse. This formulation allows the BOD module to seamlessly transition between domains where uncertainty is driven by different sources.

Overall, the BOD sampling approach is enabled by application of a fully analytical dynamics model (including the J_2 perturbation for LEO [28]) which minimizes the computation cost of orbit propagation. Similar analytic models can

be employed in other orbit regimes [31, 32]. Computation cost of the algorithm increases linearly with the number of targets in the subsystem and linearly with the number of $\delta\lambda$ samples, allowing efficient scaling.

The necessity of $\delta\lambda$ sampling is also informed by a StarFOX observability analysis, in that $\delta\lambda$ tends to be most weakly observable in swarms and formation-flying scenarios [11, 14]. This is not necessarily the case for more general configurations; in constellation scenarios with stronger observability, it may instead be possible to solve for $\delta\lambda$ directly using a similar least squares approach. Nevertheless, a limitation of the BOD algorithm is that it requires an initial guess for the system state that will converge using linearized state updates. Convergence may therefore be challenging if inter-spacecraft separations are on the order of the orbit radius, or if errors in the a-priori absolute orbit estimate are greater than several kilometers. In future, observability-aware and graph-oriented approaches will be leveraged to examine constellation cases at far separations.

E. Sequential Orbit Determination

The SOD module continually refines estimates of the orbits of all spacecraft in the subsystem as well as auxiliary parameters (e.g. sensor biases, ballistic coefficients, and differential clock offsets) by seamlessly fusing measurements from all observers transmitted over the ISL. The SOD module is based on an adaptive, efficient unscented Kalman filter (UKF) [10, 14] which applies the bearing angle measurement model and numerical GVE dynamics model from Section II. As opposed to an extended Kalman filter, the UKF is able to incorporate fully nonlinear dynamics and measurement models and preserves second order moments in the probability distribution. The resulting improvements in observability are critical for enabling maneuver-free state convergence using angles-only measurements from a single observer [14]. Three additional features are included in SOD 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 to reduce the number of calls to the orbit propagator by almost a factor of two [47]. Third, a measurement assignment procedure is used to determine whether any received measurements from remote observers are in fact measurements of targets being tracked by the local ARTMS instance. This enables multi-observer UKF measurement updates for observability and navigation performance improvements. The measurement assignment procedure is outlined as follows.

SOD first attempts to determine whether any of the remote observers ARTMS is communicating with correspond to targets being tracked by the onboard ARTMS instance. Let remote observer m have inertial Cartesian state \mathbf{x}_m^I with associated covariance \mathbf{P}_m^I (as computed from its broadcast absolute orbit estimate). Let local target n have inertial Cartesian state \mathbf{x}_n^I with associated covariance \mathbf{P}_n^I (as computed from onboard orbit estimates). The Mahalanobis distance σ_{mn} [42] between the state estimates is computed as

$$\sigma_{mn} = \sqrt{(\mathbf{x}_m^I - \mathbf{x}_n^I)^\top (\mathbf{P}_m^I + \mathbf{P}_n^I)^{-1} (\mathbf{x}_m^I - \mathbf{x}_n^I)} \quad (18)$$

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) $\sigma_{mn} \leq \sigma_{\text{assign}}$, i.e. the remote observer state is similar to the local target state
- 3) $\sigma_{pn} \geq \sigma_{\text{safe}} \forall p \neq m$, i.e. there is no other remote observer state similar to the local target state
- 4) $\sigma_{mq} \geq \Sigma_{\text{safe}} \forall q \neq n$, i.e. there is no other local target state similar to the remote observer state

Identifications are removed if $\sigma_{mn} \geq \sigma_{\text{remove}}$, with $\sigma_{\text{remove}} > \sigma_{\text{safe}} > \sigma_{\text{assign}} > 0$. These conditions ensure that targets are formally identified and that measurements broadcast by remote observer 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 of the measurements received from remote observers correspond to RSOs being tracked by the onboard ARTMS instance. Let $\mathbf{y}_{\text{meas},i}^{\mathcal{V}_m}$ be bearing angle i received from a remote observer m in its VBS frame. Let $\mathbf{y}_{\text{model},j}^{\mathcal{V}_m}$ be the modeled bearing angle of locally-tracked RSO j , as would be seen by m in its VBS frame. RSO j may be a local target, or the local observer itself. In practice, computing $\mathbf{y}_{\text{model},j}^{\mathcal{V}_m}$ requires an orbit estimate for m (either broadcast by m , or from the local ARTMS instance if m was matched to a local target); an orbit estimate for j , from the local ARTMS instance; and the rotation from \mathcal{I} to the VBS frame of m , as broadcast by m . Corresponding bearing angle covariance $P_{\text{model},j}^{\mathcal{V}_m}$ is also computed via an unscented transform of the associated state covariances. The Mahalanobis distance σ_{ij} [42] between the measured and modeled angles is then

$$\sigma_{ij} = \sqrt{(\mathbf{y}_{\text{meas},i}^{\mathcal{V}_m} - \mathbf{y}_{\text{model},j}^{\mathcal{V}_m})^\top (P_{\text{model},j}^{\mathcal{V}_m})^{-1} (\mathbf{y}_{\text{meas},i}^{\mathcal{V}_m} - \mathbf{y}_{\text{model},j}^{\mathcal{V}_m})} \quad (19)$$

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

- 1) $\sigma_{ij} \leq \sigma_{\text{match}}$, i.e. the remote measurement is close to the modeled measurement
- 2) $\sigma_{kj} \geq \sigma_{\text{ambig}} \forall k \neq i$, i.e. there is no other remote measurement that matches the local modeled measurement
- 3) $\sigma_{il} \geq \sigma_{\text{ambig}} \forall l \neq j$, i.e. there is no other local modeled measurement that matches the remote measurement

Figure 9 includes conceptual illustrations of four cases of modeled and observed measurements from a remote observer which (from left to right) show all conditions satisfied and violations of Conditions 1, 2 and 3, respectively. Together, these conditions ensure that measurements are only assigned when observed and modeled measurements uniquely agree with a statistical certainty determined by the values of σ , with $\sigma_{\text{ambig}} > \sigma_{\text{match}} > 0$. Threshold values should be selected based on the expected number of targets, relative motion geometry, sensor noise, and available orbit knowledge for general swarming missions. For scenarios similar to StarFOX, the authors have found that $\sigma_{\text{assign}} = 3$, $\sigma_{\text{safe}} = 6$, $\sigma_{\text{remove}} = 10$, $\sigma_{\text{match}} = 3$ and $\sigma_{\text{ambig}} = 6$ provide robust assignment.

Similar identity checks are also needed to allow observers to perform filter maneuver updates. ID numbers are

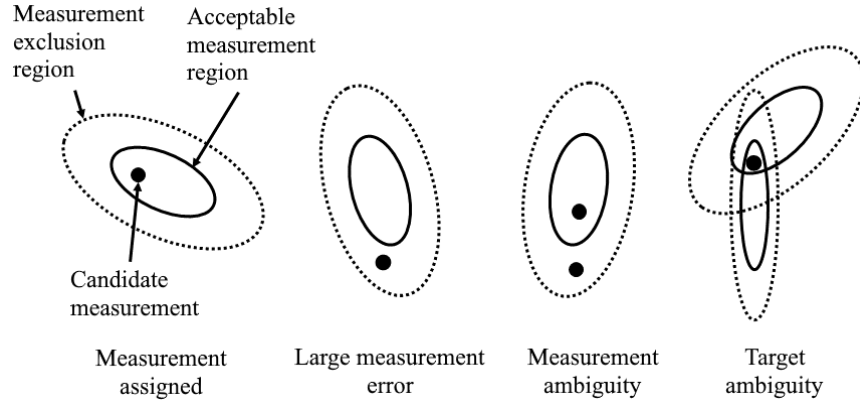


Fig. 9 Conditions in which all measurement assignment criteria are satisfied (leftmost) and conditions that violate each criteria (right).

assigned to locally-tracked objects by matching ground-provided orbit estimates to on-board orbit estimates using Mahalanobis distance thresholds. The SOD module may then use the ID numbers of provided maneuvers to perform maneuver updates for the local subsystem. Future work will investigate uncertainty-aware and multi-hypothesis methods for tracking non-cooperative targets performing potentially unknown maneuvers.

Finally, to aid robustness and error recovery, the SOD module performs autonomous health checks for the current state estimate after each filter update and attempts to automatically re-initialize if necessary. Re-initialization is performed if any of the following conditions are fulfilled:

- 1) The Mahalanobis distance between the most recent BOD relative state estimate and current SOD relative state estimate is above a user-specified threshold σ_{div} . The threshold is chosen based on expected in-orbit performance and the allowed divergence between each estimate; StarFOX applies $10 < \Sigma_{div} < 100$. The BOD and SOD estimates are only compared if the time interval between them is sufficiently small (e.g. less than 60 minutes). This ensures consistency of BOD and SOD outputs.
- 2) The number of successive filter update steps which did not feature a measurement update is above a user-specified threshold. StarFOX applies a limit of 100 measurements, or approximately 1 orbit of measurements in LEO for a VBS sample rate of 60 seconds. This ensures ARTMS has not lost sight of its targets.
- 3) The number of successive filter measurement update steps which featured a large pre-fit measurement residual is above a user-specified threshold. StarFOX applies a limit of 20 measurements to allow some time for recovery. This ensures that ARTMS has not encountered tracking errors such that its state estimates are no longer reliable.
- 4) The SOD state estimate is outside a user-specified valid range. StarFOX applies limits based on its planned orbit geometry, with $6500 \text{ km} < a < 7500 \text{ km}$; $e < 0.02$; $|\delta a| < 20 \text{ km}$; $|\delta \lambda| < 500 \text{ km}$; $|\delta e| < 20 \text{ km}$; $|\delta i| < 20 \text{ km}$. This ensures the state estimate remains physically reasonable.

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

IV. StarFOX Experiment

A. Experiment Objective

The core objective of StarFOX is to demonstrate multi-observer multi-target angles-only navigation in orbit in a more flexible, robust and autonomous fashion, building upon the contributions of prior single-observer single-target flight experiments such as ARGON [22] and AVANTI [23]. These prior experiments were affected by four key limitations: 1) reliance on accurate a-priori orbit information to initialize relative navigation for target objects (e.g. use of NORAD two-line elements), 2) reliance on access to external absolute orbit measurements to maintain state convergence for the observer (e.g. use of a GNSS receiver), 3) reliance on frequent translational maneuvers to improve angles-only range observability, and 4) use of only one observer and one target spacecraft. These limitations prevent the usage of angles-only navigation in many scenarios of interest. Specifically, reliance on accurate a-priori information and/or external orbit measurements precludes angles-only navigation in orbit regimes where accurate reference metrologies are impractical to access. Reliance on maneuvers introduces complex coupling between navigation and control systems and reduces mission lifetimes. Additionally, most DSS scenarios propose tracking of multiple targets simultaneously and/or cooperation between multiple observers. In response, StarFOX aims to explicitly demonstrate the removal of each limitation, allowing a swarm of spacecraft to perform both absolute and relative navigation using only bearing angles with minimal a-priori information. This degree of generality and autonomy enables use of angles-only navigation for a wide range of missions [31, 32, 34] and advances the readiness level of angles-only technology.

B. Starling Mission

StarFOX is a core payload of the NASA Starling swarm mission, which consists of four 6U CubeSats. Each CubeSat carries two Blue Canyon Technologies Nano Star Trackers aligned in antiparallel directions, either of which may act as the VBS for ARTMS; a GNSS receiver; an S-band radio, which acts as the ISL for ARTMS; and a cold gas propulsion system, used to conduct swarm maneuvers. Figure 10 presents an image of the four Starling spacecraft during integration at the NASA Ames Research Center. The spacecraft are denoted SN1, SN2, SN3 and SN4.

Starling will be launched into sun-synchronous low Earth orbit and nominal orbit elements of SN1 are presented in Table 1. The swarm is first deployed in an in-train (IT) or string of pearls formation mainly separated in the along-track direction with very little relative motion between spacecraft [48]. This case is the simplest operationally but is the most challenging in terms of angles-only state observability. Subsequently, the swarm will be reconfigured into two E/I-vector separated or passive safety ellipse (PSE) formations with more relative motion between spacecraft [49]. PSE formations introduce passive safety by ensuring all spacecraft are separated in the plane perpendicular to the flight direction, and furthermore possess improved observability. Nominal relative orbit elements are presented in Table 1. ARTMS performance will therefore be evaluated for three different geometric configurations, as well as during formation reconfigurations when spacecraft are maneuvering.




Fig. 10 The four Starling CubeSats during integration. Image courtesy of NASA/Dominic Hart (2022).

Figure 11 presents the relative orbits of SN2, 3, 4 in the RTN frame of SN1 for each formation. Curvilinear coordinates are used (as defined in Section II.C) to more clearly illustrate the passive safety ellipses.

Table 1 Nominal OE and ROE for the Starling swarm.

Absolute orbit						
-	a (km)	e	i ($^\circ$)	Ω ($^\circ$)	ω ($^\circ$)	M ($^\circ$)
SN1	6934	0.0020	97.5	224.8	52.3	124.9
Relative orbits (IT)						
ROE (m)	$a\delta a$	$a\delta\lambda$	$a\delta e_x$	$a\delta e_y$	$a\delta i_x$	$a\delta i_y$
SN2	0	65750	0	100	0	100
SN3	0	131500	0	100	0	100
SN4	0	197250	0	100	0	100
Relative orbits (PSE1)						
ROE (m)	$a\delta a$	$a\delta\lambda$	$a\delta e_x$	$a\delta e_y$	$a\delta i_x$	$a\delta i_y$
SN2	0	65750	0	-500	0	-500
SN3	0	131500	0	-1500	0	-1500
SN4	0	197250	0	-2000	0	-2000
Relative orbits (PSE2)						
ROE (m)	$a\delta a$	$a\delta\lambda$	$a\delta e_x$	$a\delta e_y$	$a\delta i_x$	$a\delta i_y$
SN2	0	65750	0	500	0	500
SN3	0	131500	0	-1500	0	-1500
SN4	0	197250	0	-1000	0	-1000

C. Experiment Schedule

StarFOX is divided into 14 multi-day blocks, each of which verifies a distinct form of angles-only navigation. Table 2 presents a summary of experiments planned for the initial six-month Starling mission. In the table, experiment A/B1/B2

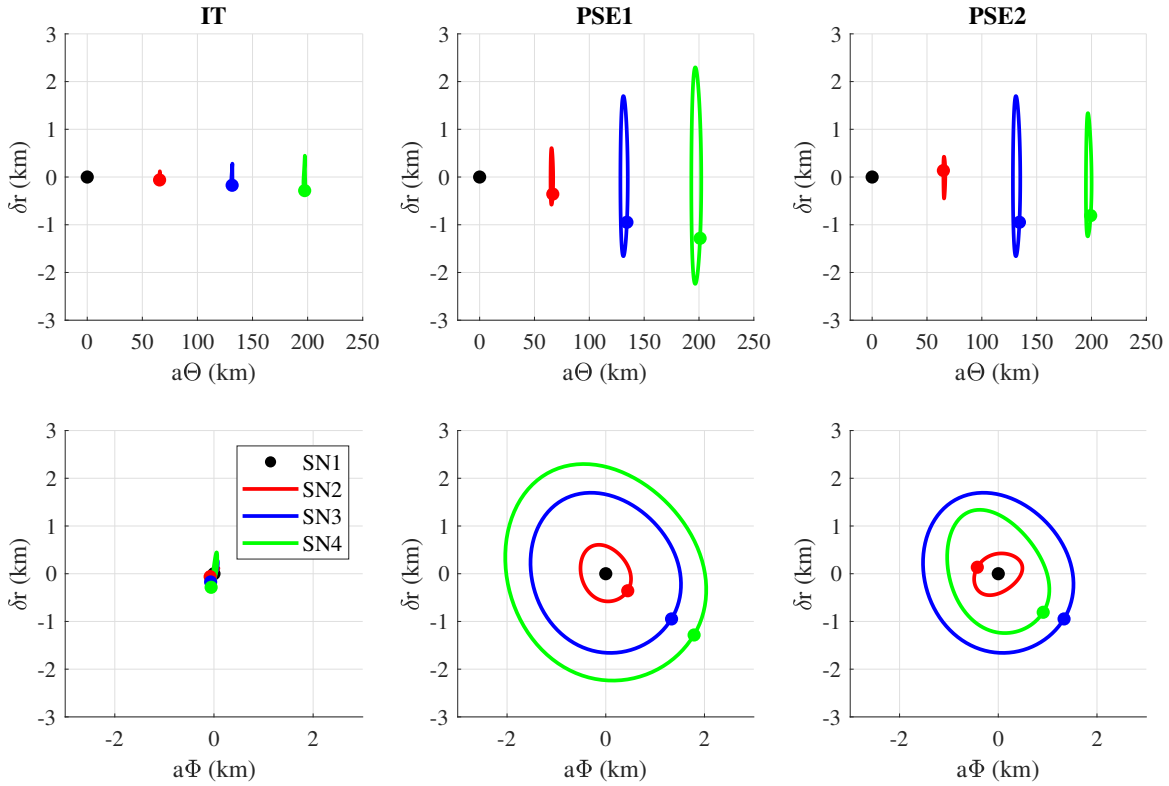


Fig. 11 Relative orbits with respect to SN1 for the IT (left), PSE1 (center) and PSE2 (right) formations.

denotes three experiments which occur simultaneously, performing navigation for 1/2/3 target spacecraft in view. The ‘attitude’ column describes the attitude of SN1/2/3/4, referring to whether its active VBS has its boresight aligned with the flight direction (+) or anti-flight direction (-). The ‘local targets’ columns list which targets are tracked by SN1/2/3/4 (left to right). The ‘remote observers’ columns list which observers SN1/2/3/4 exchanges measurements with. Remaining columns describe whether the ISL is active; whether GNSS measurements are available; whether the swarm is expected to maneuver during the experiment block; and whether an autonomous initialization from BOD is used.

Experiment blocks A/B and C/D examine single-observer multi-target navigation in an in-train formation. GNSS measurements are necessary to maintain absolute orbit convergence, and an autonomous initialization is not used due to the weaker observability of the formation, which is expected to result in large BOD state uncertainties. Blocks E/F and G/H introduce multi-observer navigation through the ISL which is expected to improve target range observability. Blocks I/J, K/L and M/N move to a PSE formation with increased observability and also begin testing the autonomous BOD initialization. Blocks O/P, Q/R and S/T display increasing autonomy by using multi-observer measurements to maintain observer absolute orbit estimates without GNSS, emulating a deep space scenario. Finally, blocks W, X and Y introduce complete distributed navigation with four cooperative observers, expected to produce the best state estimation performance as well as the highest on-board computation cost. Experiment complexity therefore increases throughout

Table 2 StarFOX experiment block configurations.

Experiment	Duration	Formation	Attitude	Local Targets				Remote Observers				ISL Active	GNSS Usage	Maneuver Present	Auto. Init.
A/B1/B2	2 days	Reconfig.	+/+/-	234	34	4	123	-	-	-	-	No	Yes	Yes	No
C/D1/D2	2 days	In-train	+/+/-	234	34	4	123	-	-	-	-	No	Yes	No	No
E/F1/F2	3 days	In-train	+/-/-	234	1	12	123	4	3	4	1	Yes	Yes	No	No
G/H1/H2	2 days	In-train	+/-/-	234	1	12	123	4	3	4	1	Yes	No	Yes	No
I/J1/J2	2 days	Reconfig.	+/+/-	234	34	4	123	-	-	-	-	No	Yes	Yes	No
K/L1/L2	2 days	PSE 1	+/+/-	234	34	4	123	4	3	4	1	No	Yes	No	Yes
M/N1/N2	3 days	PSE 1	+/-/-	234	1	12	123	4	3	4	1	Yes	Yes	No	Yes
O/P1/P2	2 days	PSE 1	+/-/-	234	1	12	123	4	3	4	1	Yes	No	No	Yes
Q/R1/R2	2 days	PSE 1	+/-/-	234	1	12	123	4	3	4	1	Yes	No	No	Yes
S/T1/T2	2 days	PSE 1	+/-/-	234	1	12	123	4	3	4	1	Yes	No	Yes	Yes
U/V1/V2	2 days	Reconfig.	+/-/-	234	1	12	123	4	3	4	1	Yes	No	Yes	Yes
W	2 days	PSE 2	+/+/-	234	34	12	123	123	134	124	123	Yes	No	Yes	Yes
X	3 days	PSE 2	+/+/-	234	34	12	123	123	134	124	123	Yes	No	No	Yes
Y	2 days	PSE 2	+/+/-	234	34	12	123	123	134	124	123	Yes	No	Yes	Yes

the schedule. Furthermore, scheduling of multiple simultaneous experiments allows performance comparisons of alternate ARTMS modes, such as different filter dynamics models, process noise handling, or measurement sample rates. Some experiment blocks include maneuvers which are known to improve angles-only observability by disambiguating target range [9]. However, maneuvers are also expected to make IMP target tracking more challenging if SOD has not yet initialized, since IMP must match maneuvers to target tracks using very limited state information.

Each experiment commences with the preparation of tables and commands to be uplinked to each spacecraft. Tables contain ARTMS software parameters, a swarm state initialization, and planned swarm maneuvers. After tables are verified on-board, each spacecraft changes its attitude to the desired configuration. ARTMS software modules are then started with nominal sample times of 60 seconds for IMP, SOD, the VBS and ISL, and 90 minutes for BOD. Experiment and bus telemetry are regularly downlinked for subsequent analysis and complete post-facto experiment playback.

V. Software-in-the-Loop Pre-Flight Verification

The ARTMS architecture is verified pre-flight through an extensive set of software simulations. Performance of each module is first validated individually, to confirm each meets the required performance standards for successful integration into the complete architecture. ARTMS is subsequently validated using Monte Carlo navigation simulations. Simulations were performed using an ARTMS flight code implementation in C++, as will be flown on board the Starling mission, and were run within a MATLAB/Simulink environment on a PC.

A. Data Generation

To generate ground truth data for test cases, the positions and velocities of swarm members were numerically integrated using the Stanford Space Rendezvous Lab’s S^3 software [50]. Dynamics models are specified in Table 3.

Spacecraft were physically modeled as 6U CubeSats with fixed attitude such that the VBS boresight points in the $\pm\hat{y}^W$ direction. In general, during StarFOX, swarm maneuvers are expected to be sufficiently large and well-known such that they act to improve angles-only observability [30]. Unless otherwise specified, maneuvers were not present in simulations to provide a more challenging observability scenario.

Five synthetic measurement types were generated from the ground truth: bearing angles, attitude quaternions, PVT navigation solutions, an a-priori swarm orbit state, and VBS images. Gaussian zero-mean measurement noise was added to measurements as per Table 4. PVT solutions, when used for either autonomous state initialization or during SOD state refinement, are assumed to possess GNSS-level accuracy (unless otherwise specified) given that StarFOX is conducted in LEO where GNSS signals are consistently available. Bearing angles and attitude quaternions possess a typical noise level for modern CubeSat star trackers [51]. Uncertainty in the a-priori state estimate is larger, to both simulate larger uncertainties after a lengthy orbit propagation period and typically larger uncertainties in the velocity direction. Bearing angle measurement gaps caused by eclipse periods and sun-blinding of the camera were modeled and occurred for between 0-60% of the orbit period. During IMP testing, 3-8 extra measurements were added to each image to emulate passing RSO and non-catalog stars, with positions pulled from a uniform distribution across the FOV.

VBS images were generated using 3D vector graphics in OpenGL [25]. The visual magnitudes of targets were modeled using a model which takes various flux sources, solar phase angle, and physical spacecraft properties into account [52]. The visual magnitudes, angles, and proper motions of stellar objects were obtained from the Hipparcos star catalog. Any objects within the camera FOV were then rendered using Gaussian point spread functions (PSF). Background noise is added to the image in the form of small variations in pixel brightness, producing centroiding errors of ~ 0.1 pixels. A small amount of noise is also added to PSF positions to emulate errors in the camera lens distortion model which is used to convert between pixel coordinates and unit vectors in \mathcal{V} . Stray light from Earth is not modeled, because the StarFOX attitude and camera have been selected to ensure that Earth is exterior to the planned FOV (and stray light was found to have minimal impact on imagery from ARGON and AVANTI [22, 23]).

The simulated camera model applies the physical properties of a Blue Canyon Technologies NST, which is the model StarFOX will use in orbit. The camera has a resolution of 1280×1024 pixels and an approximate FOV of 12×10 degrees. Images consist of greyscale pixels with brightness values from 0 to 1023.

B. Verification Performance Requirements

StarFOX mission requirements specify that the following criteria must be met by ARTMS to enable mission success:

- 1) The IMP module shall deliver measurements of target objects with combined centroiding and attitude determination errors below one pixel or $35''$ (1σ), to support BOD and SOD in producing accurate orbit solutions. This error is on par with conventional image processing algorithms and flight experience [22].
- 2) The IMP module shall assign bearing angle measurements to target objects with more than 99.5% precision. For

Table 3 Dynamics models for validation simulations.

Model	Perturbations	Propagation
IMP	Keplerian (none)	Analytic
BOD	J_2 gravity	Analytic [53]
SOD	10x10 GGM01S gravity model [2] Harris-Priester atmosphere with cannonball drag model	RK4 integrator 30s timestep
Ground truth	120x120 GGM01S gravity model [2] NRMLSISE-00 atmosphere with cannonball drag model [54] SRP with conical Earth shadow and cannonball drag model Third-body lunisolar gravity from analytic ephemeris	RK4 integrator 10s timestep

Table 4 Default noise for validation simulations.

Measurement	Noise (1σ)	Axes
Bearing angles	[20, 20]''	$[\alpha, \epsilon]$
VBS attitude	[5, 5, 20]''	$[\hat{x}^V, \hat{y}^V, \hat{z}^V]$
GNSS position	[10, 10, 10] m	$[\hat{x}^I, \hat{y}^I, \hat{z}^I]$
GNSS velocity	[0.01, 0.01, 0.01] m/s	$[\hat{x}^I, \hat{y}^I, \hat{z}^I]$
A-priori position	[100, 500, 100] m	$[\hat{x}^R, \hat{y}^R, \hat{z}^R]$
A-priori velocity	[0.1, 0.5, 0.1] m/s	$[\hat{x}^R, \hat{y}^R, \hat{z}^R]$
VBS pixel brightness	1.5%	-
VBS PSF position	[5, 5]''	$[\alpha, \epsilon]$

the fastest nominal IMP sample time of 60 seconds, this implies fewer than one false-positive measurement per target per two orbits, to support BOD and SOD in producing accurate orbit solutions.

- 3) The BOD module shall deliver an ROE solution with range error below 20% of the true inter-satellite separation when provided with two orbits of measurements. This error level is deemed sufficient to initialize SOD.
- 4) The SOD module shall demonstrate steady-state convergence within 5 orbits during maneuver-free experiments. This is on par with existing simulation results and allows sufficient time to reach convergence within each experiment block. In this context, steady-state ‘convergence’ is achieved when the rate of decrease of all ROE state uncertainties in the filter has fallen below 10% per orbit.
- 5) The SOD module shall provide estimates of the ROE state accurate to within 2% (1σ) for $\delta\lambda$ and within 1 meter per kilometer of inter-satellite separation (1σ) for all other state elements at steady-state convergence. The $\delta\lambda$ requirement stems from simulations and prior flight experiments [14, 22, 23]; other ROE requirements are derived from a worse-case measurement error budget of approximately 80'' (encompassing combined centroiding

errors, attitude determination errors, image time-tag errors and VBS mounting errors).

- 6) The ARTMS architecture shall consume no more than 50% of the on-board computer (OBC) execution time when all modules are operating at their nominal sample rates (60 seconds for IMP; 60 seconds for SOD; and 90 minutes for BOD). This ensures modules run to completion before the next call with a safety factor for unexpected computations and other software tasks. Asynchronous tasks are assumed.

C. IMP Verification

The IMP module is verified using Monte Carlo trials across a variety of absolute and relative orbits representative of the Starling mission. There are nine simulation datasets containing 100 simulations each, as defined by the magnitude of relative motion oscillations in the radial-normal plane, and the amount of noise added to PSF positions in simulated imagery. Notation is presented in Table 5. Throughout these results (and subsequent sections), ‘IMP-A’ refers to the IMP-A-1, IMP-A-2 and IMP-A-3 datasets combined; similarly, ‘IMP-1’ refers to IMP-A-1, IMP-B-1 and IMP-C-1 combined. Each simulation featured three targets uniformly distributed around three separations: [60, 70] km, [125, 135] km and [190, 200] km. Orbit elements Ω , ω , M and relative orbit phases θ , ϕ were uniformly distributed in $[0, 2\pi]$. Other OE were drawn from Starling-like ranges of $a \in [400, 700]$ km; $e \in [0.0001, 0.01]$; and $i \in [90, 105]^\circ$. Input measurements were synthetic VBS images and a single GNSS measurement at the start of the tracking period to initialize the observer absolute orbit estimate. Measurements spanned three orbits with a VBS sample time of 120 seconds.

Table 5 IMP simulation datasets.

		PSF position noise (")		
		[0, 0]	[5, 5]	[10, 10]
Magnitude of RN oscillations (m)	100	IMP-A-1	IMP-A-2	IMP-A-3
	600	IMP-B-1	IMP-B-2	IMP-B-3
	2400	IMP-C-1	IMP-C-2	IMP-C-3

Table 6 presents the mean precision, recall and accuracy of measurement assignment across each dataset as well as the mean measurement error and attitude error. Metrics are defined in terms of true positives (TP), false positives (FP), true negatives (TN) and false negatives (FN) as

$$\text{accuracy} = \frac{\text{TP} + \text{TN}}{\text{TP} + \text{TN} + \text{FP} + \text{FN}} \quad \text{precision} = \frac{\text{TP}}{\text{TP} + \text{FP}} \quad \text{recall} = \frac{\text{TP}}{\text{TP} + \text{FN}} \quad (20)$$

Accuracy assesses overall performance, precision focuses on reliability of associations, and recall focuses on frequency of associations. Precision is considered the most vital metric because angles-only orbit determination filters are sensitive to measurement errors and a single false positive can degrade the filter state estimate [38]. Here, an assignment is false positive if the assigned measurement was produced by a different target and is outside the 105" (3σ) error region

Table 6 Monte-Carlo results (1σ) for IMP simulation subsets.

Dataset	Precision (%)	Recall (%)	Accuracy (%)	Bearing Angle Error (")	Quaternion Error (")
IMP-A	99.90 \pm 0.27	85.84 \pm 10.70	96.24 \pm 2.68	9.6 \pm 6.9	10.7 \pm 6.9
IMP-B	99.60 \pm 1.23	82.89 \pm 10.43	94.31 \pm 4.69	12.1 \pm 16.6	10.4 \pm 7.1
IMP-C	99.73 \pm 0.60	84.01 \pm 11.20	95.43 \pm 3.50	11.2 \pm 11.0	10.8 \pm 7.0
IMP-1	99.81 \pm 0.39	86.57 \pm 10.21	94.96 \pm 4.14	1.8 \pm 1.0	2.4 \pm 0.3
IMP-2	99.75 \pm 0.54	83.62 \pm 10.77	94.99 \pm 4.12	11.3 \pm 4.7	10.9 \pm 1.7
IMP-3	99.71 \pm 0.61	80.05 \pm 10.72	96.03 \pm 2.92	19.7 \pm 8.3	18.5 \pm 3.6
IMP	99.78 \pm 0.49	84.23 \pm 10.81	95.24 \pm 4.08	10.3 \pm 8.8	10.6 \pm 7.0
IMP-MAN	99.49 \pm 2.03	79.10 \pm 13.35	93.97 \pm 8.71	18.1 \pm 48.3	10.6 \pm 7.0

surrounding the true target measurement.

For the IMP-A, -B and -C datasets, measurement assignment precision remains above the 99.5% threshold and false positives are minimized as desired. Despite an emphasis on discarding ambiguous measurements, recall remains above 80% and sufficient data is retained for navigation. Performance is best for IMP-A because targets possess very little motion and remain separated in the image plane. Performance is lowest for IMP-B because this scenario is most likely to produce target PSF which overlap in the FOV: sufficient relative motion is present for target tracks to potentially intersect, and motion is constrained enough to ensure targets remain in close proximity (see Figure 12). Traditional centroiding algorithms cannot distinguish this case and consider the joined PSF to be one measurement, resulting in one ‘missing’ measurement and one inaccurate measurement that is the average of the two. High overlap rates between targets are thus detrimental and formation geometry for missions such as Starling must be designed with this in mind. Nevertheless, IMP is able to continue target tracking if only occasional intersections are present. Tracking is also robust to added noise, as per the IMP-1, -2 and -3 datasets. In general, as more noise is added, precision and recall are reduced as expected, and mean bearing angle errors and attitude determination errors increase proportionally with noise.

The final IMP-MAN dataset investigates performance when one swarm spacecraft maneuvers during tracking. Maneuvers were drawn from a uniform distribution across spacecraft, execution time, direction and magnitude, with magnitude $\Delta v \in [0.1, 1]$ m/s. Maneuvers were impulsive and were executed with 5% (3σ) magnitude error and 1° (3σ) direction error. As per Table 6, the majority of maneuver cases were successfully tracked, though maneuver presence does impact reliability because resulting changes to target trajectories induce short periods of high uncertainty. However, high precision is maintained and IMP is thus expected to fulfil Starling requirements.

Optical measurement availability also has an effect on performance, noting that for StarFOX, the primary phenomena are eclipse and sun-blinding periods. Figure 13 plots data association precision versus increasing measurement gaps. Decreased measurement availability does produce lower precision because the SAMUS algorithm encounters more discontinuities during tracking and must use less information when assessing hypotheses. Nevertheless, precision

remains above the mission requirement threshold for measurement gaps of up to 60%, indicating that the presented algorithms are capable of tracking in a variety of mission scenarios.




Fig. 12 Example of connected PSF for three targets in proximity in the FOV.

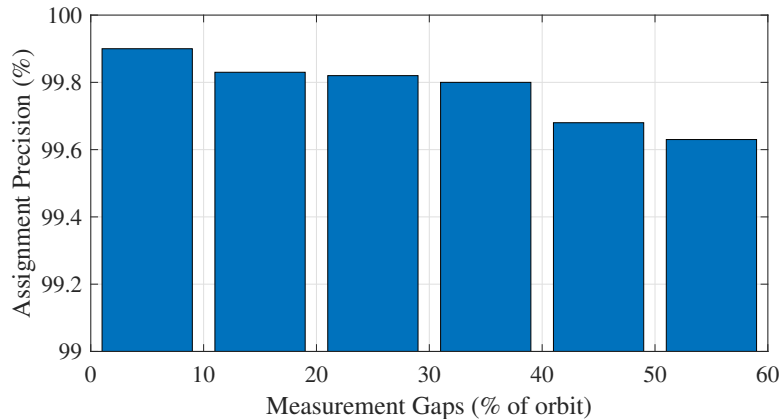


Fig. 13 Effects of measurement availability on IMP performance.

D. BOD Verification

The BOD module is verified using Monte Carlo trials across a variety of absolute and relative orbits. There are nine simulation datasets containing 100 simulations each, as defined by altitude and the magnitude of relative motion oscillations in the radial-normal plane. Notation is presented in Table 7. Each simulation featured three targets uniformly distributed around three separations: [60, 70] km, [125, 135] km and [190, 200] km. Orbit elements Ω , ω , M and relative orbit phases θ , ϕ were uniformly distributed in $[0, 2\pi]$, while other orbit properties such as inclination and eccentricity were fixed to emulate the Starling case. Input measurements were synthetic bearing angles, attitude quaternions, and a single GNSS measurement at the BOD estimation epoch to initialize the observer absolute orbit estimate. Measurements spanned two orbits with an IMP sample rate of 60 seconds.

Figure 14 presents cumulative distribution functions (CDF) for state estimation errors and 1σ uncertainties computed by the BOD algorithm across sets BOD-A, -B and -C. In the tested scenarios, altitude had minimal impact on performance

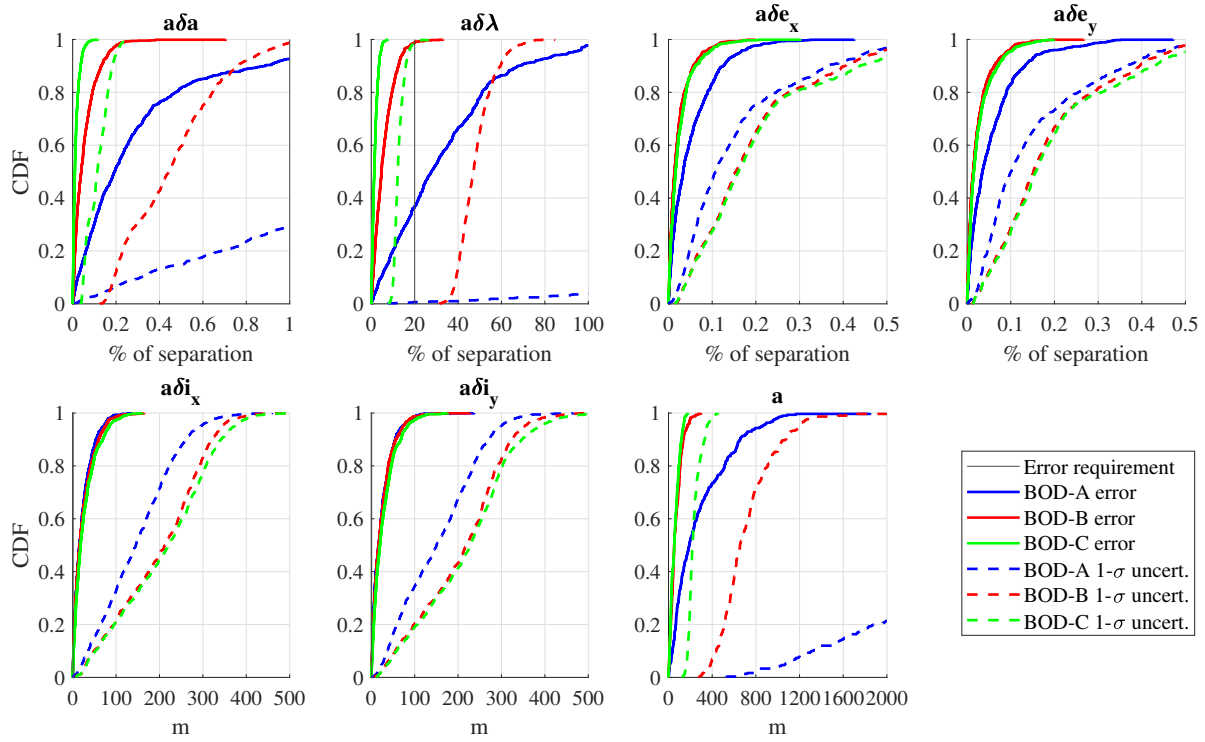


Fig. 14 CDFs for ROE state estimate errors and uncertainties during BOD simulations.

Table 7 BOD simulation datasets.

		Altitude (km)		
		400	550	700
Magnitude of RN oscillations (m)	100	BOD-A-1	BOD-A-2	BOD-A-3
	600	BOD-B-1	BOD-B-2	BOD-B-3
	2400	BOD-C-1	BOD-C-2	BOD-C-3

due to the limited effects of atmospheric drag. Thus, results are grouped by the magnitude of relative oscillations. Some estimation errors and uncertainties are also normalized by $\delta\lambda$ for convenience of comparison. Figure 15 illustrates that estimation performance for δa , $\delta\lambda$, δe_x , δe_y shows linear behavior with separation – especially in scenarios with limited relative motion – because they define the in-plane relative orbit. Performance for δi_x and δi_y , which define the out-of-plane relative orbit, displays less dependence on separation.

In the BOD-A case which emulates the Starling in-train formation, it is clear that estimates for δa and $\delta\lambda$ possess large uncertainties and errors. The $\delta\lambda$ estimate is especially poor with errors frequently on the order of target range and uncertainties much larger than target range. Semimajor axis estimation is similarly poor with errors of several hundred meters and larger uncertainties. This is because $\delta\lambda$ primarily manifests in measurements as a small bias in the elevation angle of the target. If there is very little relative motion (as in the BOD-A dataset), as well as significant

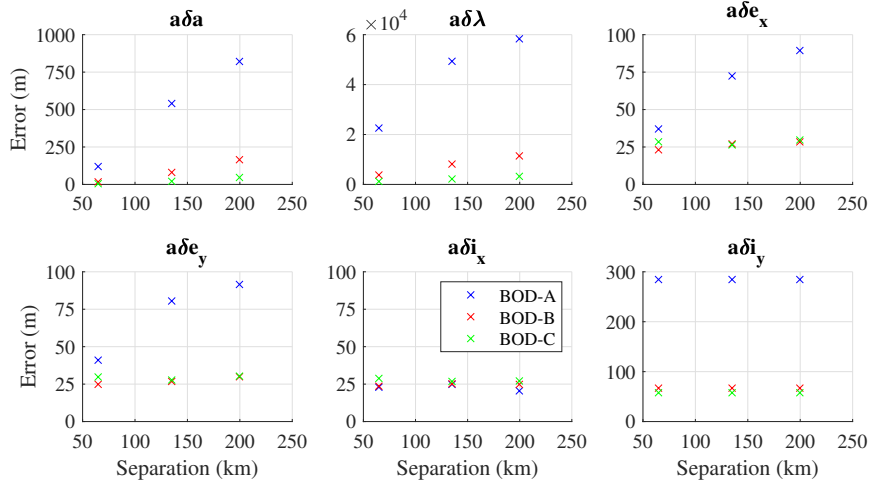


Fig. 15 Mean estimation errors vs mean separation for each target in the BOD dataset.

sensor noise, it is particularly difficult for the algorithm to determine whether its $\delta\lambda$ estimate is correct with any certainty. This subsequently leads to difficulties in distinguishing a and δa , which also affect radial position and elevation angle. Combined, $\delta\lambda$, δa and a present a weakly-observable mode. As a result, StarFOX does not intend to apply the autonomous BOD initialization during in-train operations (as per Table 2), instead choosing to rely on a-priori state information uplinked from the ground.

Estimation performance is much improved in the BOD-B case, which emulates a Starling PSE formation with moderate relative motion. Errors in $\delta\lambda$ are less than 20% of target separation in 98% of cases, whereas uncertainties are less than 50% of target separation in the majority of cases. Errors in δa are less than 0.1% of separation in 85% of cases, whereas other ROE meet this threshold in 95% of cases. This level of performance consistently approaches BOD mission performance requirements.

Finally, BOD-C case presents a PSE formation with significant relative motion. Estimation performance improves further for a , δa and $\delta\lambda$, with $\delta\lambda$ observing maximum errors of 10% of target range and maximum uncertainties of 20% of target range. As expected, larger magnitudes of oscillation in the RN plane greatly improve observability, especially for weakly-observable state components such as target range. Conversely, the estimation of more strongly-observable state components such as the relative inclination vector is more agnostic to formation geometry.

Performance is also affected by measurement availability. Figure 16 presents variations in BOD performance as measurement availability is reduced. Fewer measurements tends to lead to larger estimation errors, because the iterative least squares problem and corresponding optimization problem within BOD must be solved using less information about the target's relative orbit trajectory. Nevertheless, when considering the experiment phases and geometries during which it is intended to be used, BOD fulfils performance requirements on $\delta\lambda$ estimation in more than 98% of simulations.




Fig. 16 BOD range estimation error versus measurement availability.

E. SOD Verification

The SOD module is verified using single-observer Monte Carlo trials across a variety of absolute and relative orbits. There are nine simulation datasets containing 100 simulations each, as defined by initial position uncertainty and the magnitude of relative motion oscillations in the radial-normal plane. Notation is presented in Table 8. Orbit scenarios were similar to the prior BOD datasets, but with altitude uniformly distributed in a Starling-like range of [540, 560] km. Input measurements were synthetic bearing angles (i.e. not extracted from images), attitude quaternions, and GNSS measurements, with SOD, VBS and GNSS sample times of 60 seconds. Note that because only a single observer is present, regular GNSS updates are necessary to maintain absolute orbit convergence. The estimation period consisted of five orbits using the SOD dynamics model of Table 3. Impulsive maneuvers were also included for a specified set of trials. Maneuvers were classified as either a formation reconfiguration, in which the observer performs a large maneuver within the first two hours of the simulation with $\Delta v \in [0.1, 1]$ m/s; or formation keeping, in which randomly-chosen system members perform a small maneuver every two hours throughout the simulation with $\Delta v \in [0.01, 0.1]$ m/s. Maneuvers were drawn from a uniform distribution across execution time, magnitude and direction, and were executed with varying magnitude and direction errors (detailed later in Table 9).

Table 8 SOD simulation datasets.

		Initial position uncertainty (km)		
		5	10	20
Magnitude of RN oscillations (m)	100	SOD-A-1	SOD-A-2	SOD-A-3
	600	SOD-B-1	SOD-B-2	SOD-B-3
	2400	SOD-C-1	SOD-C-2	SOD-C-3

Figure 17a presents SOD estimation errors and 3σ uncertainties for a single trial in dataset SOD-B, along with the Monte Carlo values for SOD estimation errors and 3σ uncertainties across all trials in SOD-B. Errors are computed as the difference between the filter state estimate and ground truth simulation states. The figure also reports SOD errors and 1σ state uncertainties at the 5-orbit convergence time requirement. It is evident that SOD is able to quickly reduce

errors and uncertainties in the estimated ROE, even in the presence of significant measurement gaps. The $\delta\lambda$ estimate converges slowest, as expected, but still displays convergence without requiring maneuvers. Filter errors remain within 3σ bounds for both the individual trial and the Monte Carlo averages, demonstrating that the filter has been tuned adequately and is in good health. Figure 17b presents performance for a reconfiguration maneuver case. The impulsive maneuver after 2 hours produces a rapid decrease in the estimated ROE uncertainty, which indicates that maneuvers provide strong benefits for angles-only observability.

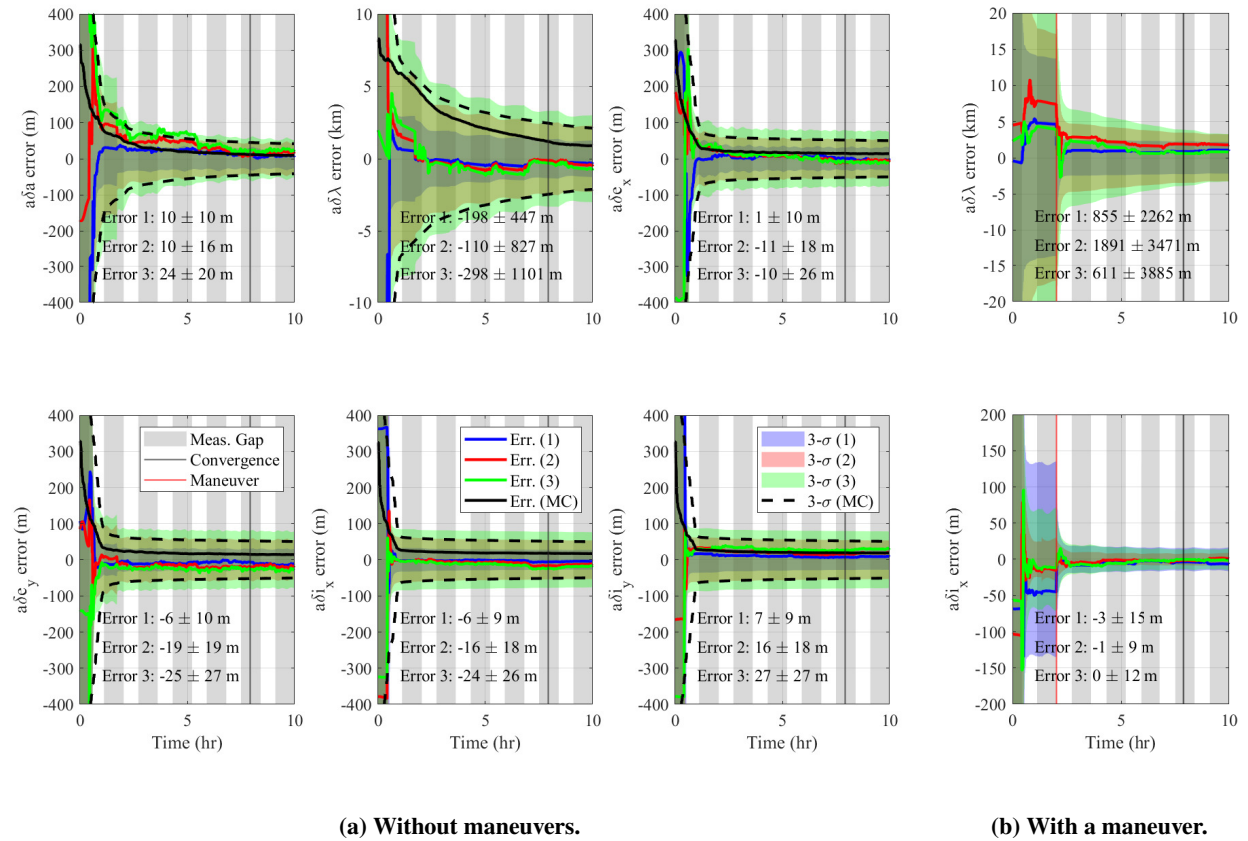


Fig. 17 SOD performance during nominal single trials and across Monte Carlo trials.

Figure 18 presents SOD state estimate uncertainty (1σ) for the 100 simulations and 300 targets in SOD-B-2, during the entire 24-hour simulation period. The Monte Carlo performance mean and its 3σ bounds are also plotted. Filter performance appears similar across all trials (with some variation due to target range, measurement gaps, and initial state errors) and state estimates remain convergent during the simulation period. Final errors are less than 1500 m in $\delta\lambda$ and less than 40 m in other ROE. Uncertainties in δa and $\delta\lambda$ decrease most significantly during the first three orbits and continue to improve over subsequent orbits. Uncertainties in the other ROE reach steady state after one orbit.

Figure 19 presents CDFs for SOD state estimation errors at the five-orbit convergence time requirement. Errors are normalized by $\delta\lambda$ for convenience of comparison. First, consider estimation of $\delta\lambda$. For SOD-1 (the best initial




Fig. 18 Standard deviations of relative state estimates for SOD operating on the SOD-B-2 dataset.

position uncertainty), errors remain below 2% in 82% of cases for SOD-A and are always below 2% for SOD-C. For SOD-3 (the worst initial position uncertainty), $\delta\lambda$ estimation errors meet the 2% specifications for 58% of SOD-A cases, 84% of SOD-B cases, and 93% of SOD-C cases. The 1σ initialization error of 20km in SOD-3 is particularly large when compared to the 65 km range of the closest target and some degradation is therefore expected. For the most challenging case of SOD-A-3, samples meeting the performance requirement tended to be those with lower initial errors; nevertheless, 82% of SOD-A-3 samples saw a reduction in $\delta\lambda$ error over the course of the simulation.

Next, consider estimation of other ROE. For SOD-1 (the best initial position uncertainty), estimation errors are below the 0.1% target range requirement for all cases studied, displaying strong steady-state performance. For SOD-3 (the worst initial position uncertainty), the most significant degradation is seen in the estimation of δa for in-train formations, such that 76% of these trials remained below the 0.1% error requirement for δa . Overall, even when the initialization is poor – due to either challenging BOD conditions or poor external state information – SOD robustly fulfils performance goals in the majority of scenarios. Figure 19 also provides a CDF for the time taken for the filter to meet steady-state error requirements. The expected trends are observed, in that this time is lengthened when initialization errors are larger and/or when the relative motion of the formation is less observable. This is most evident for the SOD-A formation, which may only meet the 5-hour convergence time requirement in 50% of cases. Measurement availability was also found to affect convergence time, primarily in less-observable scenarios, which may take one or two additional orbits to reach convergence in the presence of long eclipse or camera-blinding periods.

Finally, it is necessary to assess performance in the presence of maneuvers. Table 9 presents the influence of maneuvers




Fig. 19 CDFs for steady-state estimation errors and convergence times during SOD simulations.

on steady-state position error, uncertainty and convergence time. Simulations demonstrate that reconfigurations early in the tracking period significantly speed up convergence, especially for the in-train SOD-A scenario. Steady-state errors and uncertainties are also reduced. This is because the dynamical effects of known maneuvers on measurements allow the filter to quickly and accurately disambiguate target range. Equivalently, as maneuver execution errors increase, performance tends to worsen. Notably, a single reconfiguration is less beneficial for the SOD-B and -C scenarios (which are already more observable than SOD-A due to their relative motion). Consistent formation-keeping maneuvers in these scenarios prove more beneficial because they allow the system to improve observability and estimation performance beneath the limits provided by standard relative motion. Formation-keeping therefore implies longer convergence times also. Overall, SOD meets stated performance requirements in more than 95% of test cases, with and without swarm maneuvers, provided that state initialization errors are reasonable when compared to target range.

F. ARTMS Verification

Verification of the integrated architecture was conducted using a MATLAB Simulink model featuring four simultaneously operating ARTMS instances, to simulate the complete Starling swarm with a realistic ISL. In addition to the noise sources described in Table 4, time-tag noise of up to 20 milliseconds was added to VBS images and time-tag noise of up to 1 millisecond was added to GNSS measurements. This corresponds to the expected worst-case timing

Table 9 SOD performance with maneuvers present.

Dataset	Maneuver case	Mag. error (3σ) (%)	Dir. error (3σ) ($^\circ$)	Mean steady-state performance		
				Pos. error (1σ) (km)	Pos. uncertainty (1σ) (km)	Convergence time (hr)
SOD-A-2	None	-	-	1.55	1.78	5.48
SOD-A-2	Reconfiguration	2	1	0.68	1.19	3.00
SOD-A-2	Reconfiguration	5	2	0.71	1.20	2.98
SOD-A-2	Reconfiguration	10	5	0.80	1.27	2.95
SOD-A-2	Formation-keeping	2	1	1.17	1.27	6.26
SOD-A-2	Formation-keeping	5	2	1.19	1.27	6.26
SOD-A-2	Formation-keeping	10	5	1.28	1.27	6.26
SOD-B-2	None	-	-	0.74	1.35	3.60
SOD-B-2	Reconfiguration	5	2	0.77	1.10	2.52
SOD-B-2	Formation-keeping	5	2	0.50	0.79	4.97
SOD-C-2	None	-	-	0.46	0.75	2.31
SOD-C-2	Reconfiguration	5	2	0.83	0.74	1.93
SOD-C-2	Formation-keeping	5	2	0.36	0.32	2.78

errors when transferring data from the spacecraft bus to ARTMS during the mission. Measurement time-tags were offset from IMP, SOD and BOD sample times to enforce additional state propagation and synchronization within ARTMS.

As with individual module tests, ARTMS was validated across a variety of randomly-generated scenarios, drawn from the orbit ranges in Table 10. Performance was also investigated for the different experimental configurations in Table 2. Recall that different configurations affect swarm geometry; usage of external states or on-board methods for initialization of navigation; and whether GNSS measurements or only bearing angles are used for observer absolute orbit estimation. Ten four-spacecraft simulations were performed for each configuration. Each simulation generates four absolute orbit estimates (one per observer) and nine or ten relative orbit estimates, depending on the attitude configuration. The simulation period consisted of 24 hours with samples rates of 120 seconds for synthetic VBS images, GNSS, IMP and SOD and 3 hours for BOD. Adaptive process noise estimation was applied in SOD.

Table 10 OE and ROE ranges for ARTMS simulations.

Observer OE	Range	Target ROE	Range
a (km)	[6883, 6983]	δa (km)	[-0.05, 0.05]
e	[0.0001, 0.01]	$\delta \lambda$ (km)	[30, 300]
i ($^\circ$)	[93, 103]	δe_x (km)	[-5, 5]
Ω ($^\circ$)	[0, 360]	δe_y (km)	[-5, 5]
ω ($^\circ$)	[0, 360]	δi_x (km)	[-5, 5]
M ($^\circ$)	[0, 360]	δi_y (km)	[-5, 5]

Figure 20 presents CDFs for relative orbit estimation errors in four configurations:

- K/L: 1 observer, PSE formation, BOD initialization, using GNSS for absolute orbit updates




Fig. 20 CDFs for ROE estimation errors 24 hours after ARTMS initialization.

- C/D: 1 observer, IT formation, ground initialization, using GNSS for absolute orbit updates
- E/F: 2 observers, IT formation, ground initialization, using GNSS for absolute orbit updates
- O/P: 2 observers, PSE formation, BOD initialization, using only bearing angles for absolute orbit updates

PSE formations meet ARTMS performance goals in all trials, whereas IT formations meet performance requirements in 88% of trials (single observer) and 96% of trials (two cooperative observers). The $\delta\lambda$ error requirement proves the most challenging to fulfill. However, multi-observer measurement fusion is able to significantly reduce state estimate error, demonstrating the observability advantage gained by a distributed stereo-vision approach.

Figure 21 presents CDFs for absolute orbit estimation errors in four configurations:

- G/H: 2 observers, IT formation, ground initialization, using GNSS for absolute orbit updates
- O/P: 2 observers, PSE formation, BOD initialization, using only bearing angles for absolute orbit updates
- X-IT: 4 observers, IT formation, ground initialization, using only bearing angles for absolute orbit updates
- X: 4 observers, PSE formation, BOD initialization, using only bearing angles for absolute orbit updates

Orbit element errors are multiplied by the observer semimajor axis for interpretation as position errors in meters. Although no specific performance goals are put forth for absolute orbit estimation using angles-only measurements, it can be observed that PSE formations produce smaller estimation errors, as expected. The semimajor axis is most observable (further justifying its estimation within BOD also) whereas the mean argument of latitude is least observable. There exists a poorly observable mode whereupon errors in u are closely correlated to errors in $\delta\lambda$, since both state components affect spacecraft positioning along the orbit path. Absolute position errors after 24 hours are typically

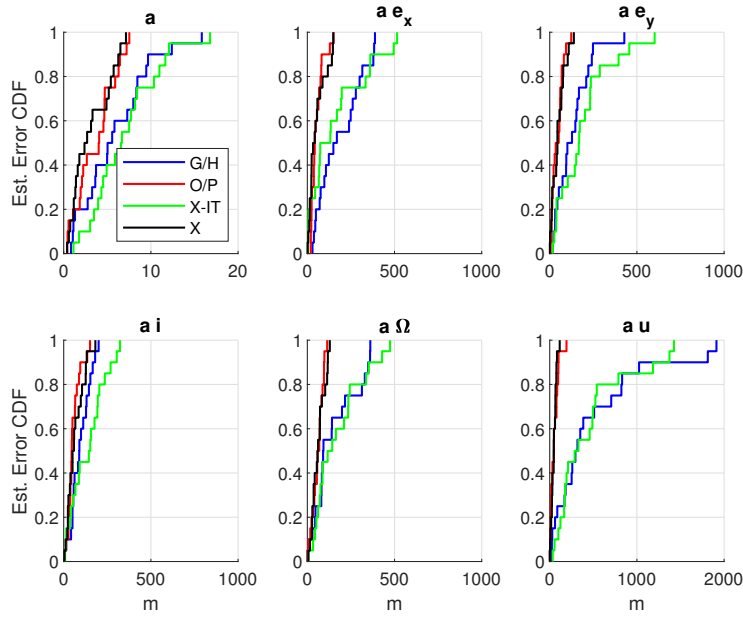


Fig. 21 CDFs for OE estimation errors 24 hours after ARTMS initialization.

on the order of hundreds of meters and it is therefore expected that observability will be sufficient for ARTMS to demonstrate angles-only absolute orbit determination in flight.

Table 11 presents a summary of mean position estimation errors and uncertainties in different configurations. Some trials apply GNSS corruption, in which GNSS measurements possess 10 times the default amount of noise, to emulate scenarios in which absolute position information is of lower quality (e.g. deep space scenarios). Some trials also include known impulsive maneuvers by the observer with $\Delta v \in [0.1, 1]$ m/s. Maneuvers possess random execution times, magnitudes and directions and were executed with 5% (3σ) magnitude error and 1° (3σ) direction error.

In Table 11, the ARTMS architecture meets performance requirements in that relative position estimates and uncertainties are consistently below 2% of target range. Maneuvers continue to aid observability, most obviously for the single-observer in-train case (configuration C/D) which possesses the weakest observability. Maneuvers have a much smaller impact in multi-observer cases (configuration O/P) which may already disambiguate range via measurement sharing. In case K/L (which does receive regular GNSS measurements), GNSS corruption produces the expected increase in absolute state errors but has little effect on relative state errors. However, in case O/P (which does not receive regular GNSS measurements), GNSS corruption reduces the quality of the one-off absolute orbit estimate used in BOD. This in turn reduces the quality of the output BOD swarm state estimate, slightly worsening absolute and relative estimation performance. However, the performance degradation is minor and ARTMS is overall robust to GNSS errors. Adaptive process noise estimation is also shown to reduce state uncertainty, supporting its usage in orbit.

Table 11 Mean position estimation errors and covariances for ARTMS simulations.

Config.	Modification	Absolute position		Relative position	
		Error (1σ)	Uncertainty (1σ)	Error (1σ)	Uncertainty (1σ)
C/D	-	5 ± 2 m	12 ± 0 m	$0.72 \pm 0.78\%$	$0.70 \pm 0.50\%$
C/D	PSE formation	6 ± 3 m	12 ± 0 m	$0.10 \pm 0.10\%$	$0.21 \pm 0.20\%$
C/D	Maneuvers present	5 ± 2 m	12 ± 1 m	$0.44 \pm 0.41\%$	$0.50 \pm 0.31\%$
E/F	-	7 ± 3 m	13 ± 2 m	$0.21 \pm 0.36\%$	$0.19 \pm 0.24\%$
G/H	-	670 ± 480 m	900 ± 400 m	$0.58 \pm 0.92\%$	$0.68 \pm 0.60\%$
K/L	-	6 ± 2 m	13 ± 1 m	$0.10 \pm 0.15\%$	$0.28 \pm 0.26\%$
K/L	No adaptive proc. noise	6 ± 2 m	20 ± 0 m	$0.12 \pm 0.18\%$	$0.42 \pm 0.33\%$
K/L	GNSS corruption	21 ± 8 m	59 ± 2 m	$0.11 \pm 0.13\%$	$0.29 \pm 0.26\%$
O/P	-	140 ± 40 m	360 ± 200 m	$0.05 \pm 0.06\%$	$0.13 \pm 0.10\%$
O/P	No adaptive proc. noise	170 ± 40 m	480 ± 140 m	$0.04 \pm 0.05\%$	$0.14 \pm 0.11\%$
O/P	Maneuvers present	130 ± 30 m	330 ± 180 m	$0.05 \pm 0.05\%$	$0.12 \pm 0.08\%$
O/P	GNSS corruption	230 ± 80 m	370 ± 210 m	$0.05 \pm 0.06\%$	$0.14 \pm 0.11\%$
X	-	150 ± 40 m	250 ± 120 m	$0.05 \pm 0.05\%$	$0.12 \pm 0.10\%$
X	In-train formation	670 ± 350 m	690 ± 240 m	$0.41 \pm 0.46\%$	$0.53 \pm 0.41\%$

VI. Hardware-in-the-Loop Pre-Flight Verification

ARTMS is next verified using hardware in the loop elements. First, navigation simulations featuring a Blue Canyon Technologies NST verify that ARTMS can operate on real camera images. Second, night sky image processing tests featuring the integrated star trackers on board each spacecraft verify that ARTMS can obtain accurate measurements from integrated cameras. Third, navigation simulations running on a CubeSat OBC verify that ARTMS execution times are compliant when run on CubeSat hardware.

A. Camera-in-the-Loop Simulations

HIL images were retrieved from a Blue Canyon Technologies NST as stimulated by the Stanford SLAB Optical Simulator (OS). The OS is a variable-magnification testbed consisting of two lenses and a microdisplay. Synthetic space scenes are generated and shown on the display and by moving the lenses and display relative to each other, the VBS under test is stimulated with appropriate magnification. The OS is calibrated such that the VBS image is similar in radiosity and geometry to what would be observed in orbit. Development, calibration, and usage of the OS are detailed in [25] with achievable errors between desired and measured bearing angles of less than $10''$. HIL images were taken for two IT scenarios and two PSE scenarios across 16 hours and subsequently ingested into the MATLAB/C++ simulation framework. Other measurement inputs were the same as during software-only simulations.

Figures 22 and 23 present SOD estimation errors and uncertainties from a camera-in-the-loop simulation of experiment O/P. Absolute orbit statistics are multiplied by the semimajor axis for interpretation as position errors. During the simulation, the IMP module provides consistent measurement assignment and the BOD initialization produces initial

state errors on the order of several kilometers. The SOD absolute and relative state estimates subsequently converge to steady state values within several hours without maneuvers, and measurement sharing is utilized to achieve improved relative orbit determination and angles-only absolute orbit determination. Table 12 compares performance when using

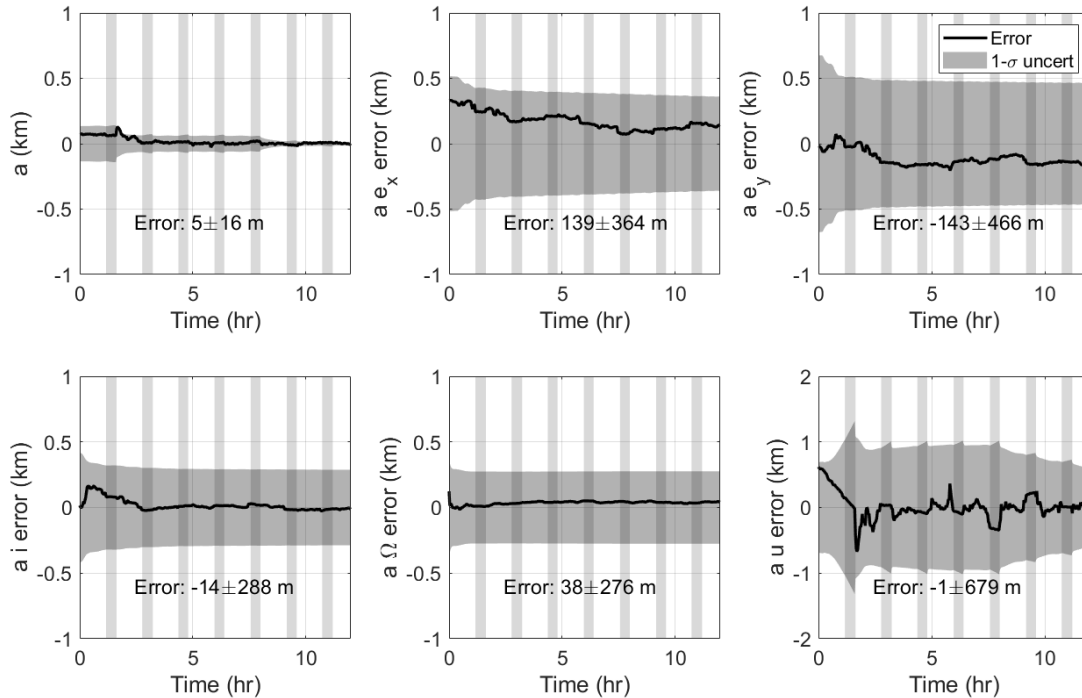


Fig. 22 Estimation errors and formal covariances for the observer absolute orbit estimate.

synthetic imagery versus HIL imagery for four experiment configurations. Results are comparable between synthetic

Table 12 State estimation performance comparison for HIL simulations.

Config.	Modification	Position error (1σ)	
		Absolute	Relative
C/D	-	5 ± 16 m	$0.82 \pm 1.26\%$
C/D	HIL	5 ± 16 m	$0.71 \pm 1.18\%$
K/L	-	5 ± 18 m	$0.13 \pm 0.43\%$
K/L	HIL	5 ± 18 m	$0.15 \pm 0.40\%$
X	PSE	90 ± 360 m	$0.07 \pm 0.21\%$
X	HIL, PSE	160 ± 360 m	$0.05 \pm 0.20\%$
X	In-train	520 ± 1100 m	$0.44 \pm 0.97\%$
X	HIL, in-train	630 ± 1060 m	$0.21 \pm 0.98\%$

and HIL trials, indicating that ARTMS is able to operate on real camera imagery.

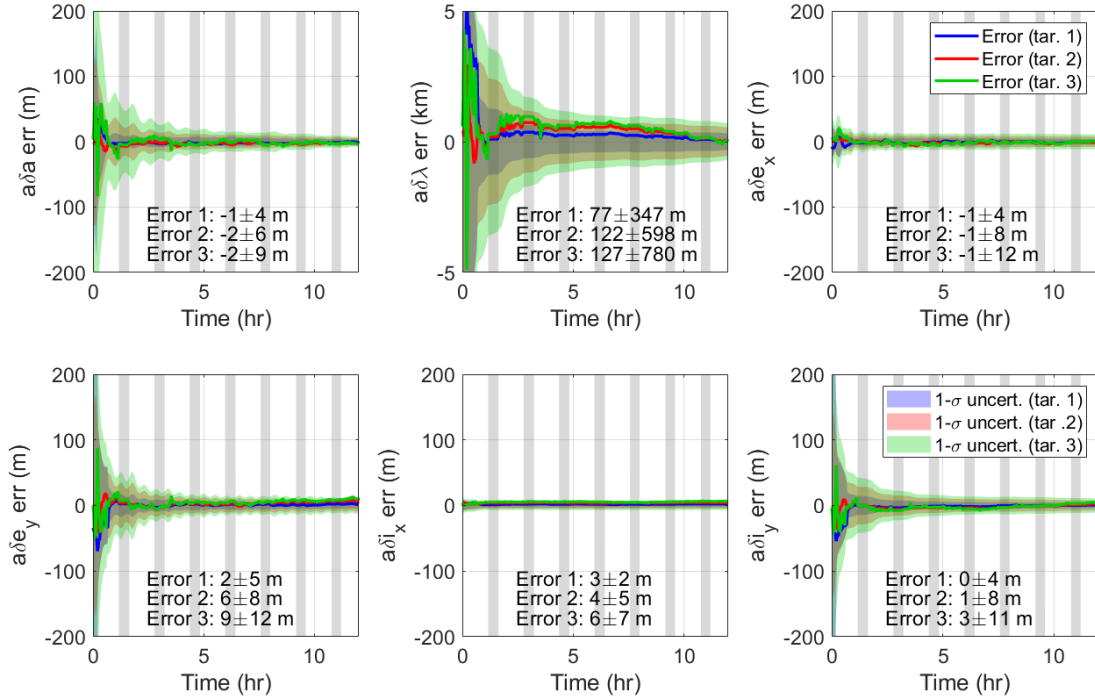


Fig. 23 Estimation errors and formal covariances for target relative orbit estimates.

B. Integrated Camera Night Sky Tests

The level of bearing angle measurement error when using integrated hardware is investigated by comparing the inertial unit vector measurements produced by star tracker on-board algorithms (previously verified by the manufacturer) to the inertial unit vector measurements produced by ARTMS algorithms. Images containing point sources with known locations are obtained by taking night sky images with each Starling spacecraft, as in Figure 24. ‘Ground truth’ inertial unit vectors are those of identified stars as obtained from the Hipparcos star catalog, expressed in \mathcal{I} . Measured ARTMS unit vectors to identified stars in \mathcal{V} are obtained using the IMP algorithm described in Section III.C. Vectors are then rotated into \mathcal{I} for comparison using the attitude estimate ${}^{\mathcal{V}}\overline{\mathbf{R}}^{\mathcal{I}}$ computed by ARTMS. A key step is the conversion of pixel coordinates to true unit vectors, which requires use of a camera model to compensate for known lens distortions. Distortions may introduce unit vector errors on the order of tens or hundreds of arcseconds, preventing successful star identification and introducing significant measurement errors. The distortion model used by ARTMS is polynomial with calibrated coefficients for each individual camera. Note that this model was not used in Section VI.A because camera distortions became coupled with the OS facility distortions and were thus removed during facility calibration.

Table 13 presents inertial unit vector errors. Errors remain well below the stated IMP performance requirement of $35''$ (1σ), indicating that the camera distortion model is appropriate and that the star trackers and IMP are able to provide suitable angles-only measurements to BOD and SOD. The larger errors observed by SN4 Tracker 1 are caused



Fig. 24 A night sky taken by Starling SN4 from the roof of the NASA Ames Research Center.

Table 13 Measurement errors of integrated star trackers.

Spacecraft	Tracker 1 Meas. Error (")		Tracker 2 Meas. Error (")	
	[Min., Max.]	Mean (1σ)	Range	Mean (1σ)
SN1	[5.8, 26.9]	12.0 ± 4.3	[3.2, 24.6]	13.9 ± 6.3
SN2	[3.2, 20.9]	11.3 ± 4.7	[5.4, 19.2]	12.8 ± 4.3
SN3	[1.6, 16.4]	7.4 ± 3.7	[1.2, 17.0]	10.4 ± 5.0
SN4	[5.4, 72.3]	13.1 ± 8.7	[2.1, 16.6]	8.4 ± 5.1

by a poorer calibration model in the far upper left corner of the image plane. However, the model remains sufficient for star identification and furthermore, target RSO will typically be located much closer to the image center.

C. CubeSat Microprocessor Performance Profiling

ARTMS modules were profiled on a Tyvak ARM Cortex-A8+ microprocessor running at 800 MHz. For comparison, the Starling spacecraft employ Xilinx Q7S ARM Cortex-A9 microprocessors running at up to 766 MHz. During profiling, the IMP and SOD modules performed 10 successive executions on pre-generated test data for experiment configuration M/N, whereas the BOD module was executed once. The primary computation costs for SOD arise from the number of measurement updates, dynamics model fidelity, and amount of state propagation performed within the filter. Primary computation costs for IMP arise from the numbers of tracked targets and visible stars and whether a prior attitude estimate is provided to assist with star identification. Primary computation costs for BOD arise from the number of targets, number of range samples, and measurement batch size. Experiment M/N is considered typical in that it features two cooperative observers each obtaining bearing angles to three targets, with regular GNSS absolute orbit updates. To increase computation costs, no attitude estimate was provided to IMP, and additional state propagation was

enforced within SOD to synchronize measurement epochs to module execution epochs. The SOD dynamics model was that in Table 3. BOD processed a 200-measurement batch using 150 range samples, which is the maximum number of measurements with a nominal number of range samples.

Table 14 presents module execution times across 10 tests. For the fastest nominal sample times of 60 seconds (IMP

Table 14 ARTMS execution times during performance profiling.

Module	Execution time (s)		
	[Min., Max.]	Mean (1σ)	Max. % of sample time
SOD	[9.81, 12.09]	10.68 ± 0.89	20.2
IMP	[6.48, 12.61]	8.03 ± 2.45	21.0
BOD	[610.8, 611.6]	611.1 ± 0.27	11.3
ALL	-	-	52.5

and SOD) and 90 minutes (BOD), the maximum observed execution times require approximately 20% of OBC time (IMP and SOD) and 10% of OBC time (BOD). This implies that ARTMS may require ~50% of total OBC time, close to the mission requirement constraint. However, if anomalous performance is observed in orbit, ARTMS parameters may be tuned to reduce computation by e.g. reducing the number of allowed target track hypotheses in IMP, reducing SOD dynamics model fidelity, or reducing the number of BOD range hypotheses, without significant impacts on accuracy.

VII. Conclusion

This paper presents an overview of the Starling Formation-flying Optical eXperiment (StarFOX), intended as the first in-flight demonstration of autonomous distributed angles-only navigation for a spacecraft swarm. Prior demonstrations of angles-only navigation in orbit have displayed only single-observer single-target navigation and have required external orbit information and/or system maneuvers to achieve orbit determination. StarFOX instead applies the new angles-only Absolute and Relative Trajectory Measurement System (ARTMS), an architecture that provides autonomous, distributed, and scalable navigation for distributed systems in deep space using inexpensive optical sensors. Its multi-observer multi-target framework is able to achieve complete swarm orbit determination with minimal external orbit information and without maneuvers. ARTMS consists of three modules: image processing (IMP), batch orbit determination (BOD), and sequential orbit determination (SOD). The IMP module provides batches of measurements to each observed target using time-tagged images from the onboard camera and a single coarse estimate of the observer’s orbit. The BOD module uses these batches of bearing angles to each target, along with the observer orbit estimate, to provide estimates of the orbits of each observed system member. Finally, the SOD module refines system orbit estimates (as well as auxiliary parameter estimates such as ballistic coefficients or differential clock offsets) by seamlessly fusing measurements from multiple observers received over an inter-satellite link within an unscented Kalman filter.

To validate each algorithm, module performance is assessed via Monte Carlo simulations of representative StarFOX

scenarios in the context of mission performance requirements. IMP simulations demonstrate measurement assignment precision of more than 99.9% in near-circular orbits and measurement errors below one pixel (1σ). BOD simulations demonstrate target range estimation errors of less than 20% in more than 95% of trials. SOD simulations demonstrate convergence to target range errors of less than 2%, and errors in other ROE which are less than 0.1% of target range, within five orbits in more than 95% of trials. Performance is affected by the amount of relative motion in the swarm, the amount of state initialization error, and the frequency of target overlap in the field of view, but is overall robust and meets requirements in the majority of cases.

The complete ARTMS architecture is verified with software- and hardware-in-the-loop simulations of StarFOX experiments using ARTMS flight code, CubeSat star trackers and a CubeSat on-board computer. ARTMS is able to autonomously initialize navigation and demonstrates measurement sharing and fusion between multiple observers, achieving simultaneous absolute and relative orbit determination using angles-only measurements. Angles-only absolute position errors converge to less than 500 m if the swarm possesses sufficient relative motion. ARTMS is also able to operate on representative imagery collected by a star tracker and execution times on a CubeSat microprocessor remain below 50% of nominal sample times. Results verify that ARTMS is capable of fulfilling StarFOX objectives, thus increasing the readiness level of critical technologies for future deep space or space domain awareness missions.

Funding Sources

This work was supported by the NASA Small Spacecraft Technology Program (cooperative agreement number 80NSSC18M0058) for contribution to Starling/StarFOX efforts, and by the Air Force Office of Scientific Research and Dr. Michael Yakes (award number FA9550-21-1-0414) for the project titled ‘Autonomous Distributed Angles-Only Orbit Determination using Multiple Observers’.

Acknowledgments

The authors acknowledge the invaluable contributions of Joshua Sullivan towards the conception and development of StarFOX and ARTMS, and of Toby Bell towards the development of the ARTMS flight software.

References

- [1] D’Amico, S., Pavone, M., Saraf, S., Alhussien, A., Mohammed Al-Saud, T. S., 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*, 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.
- [3] Krieger, G., Moreira, A., Fiedler, H., Hajnsek, I., Werner, M., Younis, M., and Zink, M., “TanDEM-X: A Satellite Formation

- for High-Resolution SAR Interferometry,” *IEEE Transactions on Geoscience and Remote Sensing*, Vol. 45, No. 11, 2007, pp. 3317–3341.
- [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.
- [5] Burch, J., Moore, T., Torbert, R., and Giles, B., “Magnetospheric Multiscale Overview and Science Objectives,” *Space Science Reviews*, Vol. 199, No. 1-4, 2016, pp. 5–21.
- [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.
- [7] Arney, D., Sutherland, R., Mulvaney, J., Steinkoenig, D., Stockdale, C., and Farley, M., “On-orbit Servicing, Assembly, and Manufacturing (OSAM) State of Play,” Tech. rep., NASA, 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, Advanced Technologies II, SSC13-III-5*, Logan, UT, 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. URL http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=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*, San Antonio, TX, 2017.
- [11] Koenig, A. W., and D’Amico, S., “Observability-Aware Numerical Algorithm for Angles-Only Initial Relative Orbit Determination,” *2020 AAS/AIAA Astrodynamics Specialist Conference*, Lake Tahoe, CA, 2020.
- [12] Yim, J. R., Crassidis, J. L., and Junkins, J. L., “Autonomous Orbit Navigation of Two Spacecraft System Using Relative Line of Sight Vector Measurements,” *AAS, Vol. AAS, Vol. Paper 04-257*, 2004. URL <http://ai2-s2-pdfs.s3.amazonaws.com/ef48/18cfc1deaf9f3c69cf3069a42219c5180eaf.pdf>.
- [13] Greaves, J., and Scheeres, D., “Optical Only Co-Estimation and Maneuver Classification in the Cislunar Regime,” *Rocky Mountain AAS GNC Conference*, Breckenridge, Colorado, 2022.
- [14] 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.
- [15] “Distributed orbit determination and observability analysis for satellite constellations with angles-only measurements,” *Automatica*, Vol. 129, 2021, p. 109626.
- [16] 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.

- [17] Willis, M., and D'Amico, S., "Fast Angles-Only Relative Navigation Using Polynomial Dynamics," *11th International Workshop on Satellite Constellations and Formation Flight*, Milano, Italy, 2022.
- [18] Kulik, J., and Savransky, D., "State Transition Tensors for Passive Angles-Only Relative Orbit Determination," *33rd AAS/AIAA Space Flight Mechanics Meeting*, Austin, Texas, 2023.
- [19] 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, 2015.
- [20] LeGrand, K., and DeMars, K., "Relative multiple space object tracking using intensity filters," *Proceedings of the 18th International Conference on Information Fusion*, Washington, DC, 2015.
- [21] Kruger, J., and D'Amico, S., "Autonomous angles-only multitarget tracking for spacecraft swarms," *Acta Astronautica*, Vol. 189, 2021, pp. 514–529.
- [22] 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.
- [23] 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.
- [24] Sanchez, H., McIntosh, D., Cannon, H., Pires, C., Sullivan, J., D'Amico, S., and O'Connor, B., "Starling-1: Swarm Technology Demonstration," *32nd Annual Small Satellite Conference*, Logan, UT, 2018.
- [25] Beierle, C., and D'Amico, S., "Variable Magnification Optical Stimulator for Training and Validation of Spaceborne Vision-Based Navigation," *Journal of Spacecraft and Rockets*, Vol. 56, No. 4, 2019, pp. 1060–1072.
- [26] Vallado, D. A., and McClain, W. D., *Fundamentals of Astrodynamics and Applications*, 4th ed., Microcosm Press, Hawthorne, California, 2013.
- [27] D'Amico, S., "Autonomous Formation Flying in Low Earth Orbit," Ph.D. thesis, Delft University, 2010.
- [28] 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.
- [29] 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>.
- [30] 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*, Snowbird, UT, 2018.
- [31] Kruger, J., Wallace, K., Koenig, A. W., and D'Amico, S., "Autonomous Angles-Only Navigation for Spacecraft Swarms around Planetary Bodies," *2021 IEEE Aerospace Conference*, Big Sky, Montana, 2021.

- [32] Iiyama, K., Kruger, J., and D’Amico, S., “Autonomous Distributed Angles-Only Navigation and Timekeeping in Lunar Orbit,” *ION International Technical Meeting*, Long Beach, California, 2022.
- [33] Kruger, J., and D’Amico, S., “Observability Analysis and Optimization for Angles-Only Navigation of Distributed Space Systems,” *11th International Workshop on Satellite Constellations and Formation Flight*, Milano, Italy, 2022.
- [34] Kruger, J., D’Amico, S., Roscoe, C., and Westphal, J., “Angles-Only Tracking and Navigation for Approach and Rendezvous in Geosynchronous Orbits,” *33rd AAS/AIAA Space Flight Mechanics Meeting*, Austin, Texas, 2023.
- [35] Alfried, K. T. (ed.), *Spacecraft Formation Flying: Dynamics, Control, and Navigation*, Butterworth-Heinemann/Elsevier, 2010.
- [36] Montenbruck, O., and Gill, E., *Satellite Orbits: Models, Methods and Applications*, Springer, 2012.
- [37] 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*, Napa, California, 2016. URL https://people.stanford.edu/damicos/sites/default/files/aassfm2016_sullivankoenigdamico.pdf.
- [38] 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.
- [39] Delabie, T., De Schutter, J., and Vandenbussche, B., “An Accurate and Efficient Gaussian Fit Centroiding Algorithm for Star Trackers,” *The Journal of the Astronautical Sciences*, Vol. 61, 2014, pp. 60–84.
- [40] 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.
- [41] Wertz, J. R., *Spacecraft Attitude Determination and Control*, Springer Science & Business Media, 2012.
- [42] Blackman, S. S., and Popoli, R., *Design and Analysis of Modern Tracking Systems*, Artech House, Boston, Massachusetts, 1999.
- [43] Bar-Shalom, Y., Willett, P. K., and Tian, X., *Tracking and Data Fusion: A Handbook of Algorithms*, YBS Publishing, Storrs, Connecticut, 2011.
- [44] Kisantal, M., Sharma, S., Park, T. H., Izzo, D., Märtens, M., and D’Amico, S., “Satellite Pose Estimation Challenge: Dataset, Competition Design and Results,” *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 56, No. 5, 2020, pp. 4083–4098.
- [45] Ester, M., Kriegel, H.-P., Sander, J., and Xu, X., “A Density-Based Algorithm for Discovering Clusters 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*, 1996, p. 226–231.
- [46] Blackman, S., “Multiple Hypothesis Tracking For Multiple Target Tracking,” *IEEE Aerospace and Electronic Systems Magazine*, Vol. 19, No. 1, 2004, pp. 5–18.

- [47] Stacey, N., and D'Amico, S., "Autonomous Swarming for Simultaneous Navigation and Asteroid Characterization," *AAS/AIAA Astrodynamics Specialist Conference*, Snowbird, UT, 2018.
- [48] Schmidt, J., Phillips, M., and Ruschmann, M., "Satellite Cluster Flight Design Considerations," *24th International Symposium on Space Flight Dynamics*, Laurel, Maryland, 2014.
- [49] D'Amico, S., and Montenbruck, O., "Proximity Operations of Formation-Flying Spacecraft Using an Eccentricity/Inclination Vector Separation," *Journal of Guidance, Control, and Dynamics*, Vol. 29, No. 3, 2006, pp. 554–563.
- [50] Giraldo, V., and D'Amico, S., "Development of the Stanford GNSS Navigation Testbed for Distributed Space Systems," *Proceedings of the 2018 International Technical Meeting of the Institute of Navigation*, Reston, VA, 2018, pp. 837–856.
- [51] Hegel, D., "Small Spacecraft Subsystem State-of-the-Art: Attitude Determination and Control," *Proceedings of the 15th International Planetary Probe Workshop*, The University of Colorado Boulder, 2018.
- [52] Cognion, R. L., "Large phase angle observations of GEO satellites," *Sensors and Systems for Space Applications VI*, Vol. 8739, edited by K. D. Pham, J. L. Cox, R. T. Howard, and G. Chen, International Society for Optics and Photonics, SPIE, 2013, pp. 194–205.
- [53] Brouwer, D., "Solution of the problem of artificial satellite theory without drag." *Astronomical Journal*, Vol. 64, No. 1274, 1959, pp. 378–397.
- [54] Picone, J., Hedin, A., Drob, D. P., and Aikin, A., "NRLMSISE-00 Empirical Model of the Atmosphere: Statistical Comparisons and Scientific Issues," *Journal of Geophysical Research: Space Physics*, Vol. 107, No. A12, 2002, pp. SIA–15.