

Starling Formation-Flying Optical Experiment: Initial Operations and Flight Results

Justin Kruger, Simone D'Amico
Stanford University
496 Lomita Mall, Stanford, CA 94305
jjkruger@stanford.edu

Soon S. Hwang
NASA Ames Research Center
Moffett Field, CA 94035
soon.s.hwang@nasa.gov

ABSTRACT

This paper presents initial flight results for angles-only navigation of a swarm of small spacecraft, conducted during the Starling Formation-Flying Optical Experiment (StarFOX). StarFOX is one of four experiments aboard the NASA Starling mission, which consists of four CubeSats launched in July 2023. Angles-only methods apply inter-satellite bearing angles obtained by on-board cameras for navigation, increasing satellite autonomy and enabling new mission concepts. Nevertheless, prior flight demonstrations have only featured one observer and target and have relied upon a-priori target orbit knowledge for initialization, translational maneuvers to resolve target range, and external absolute orbit updates to maintain convergence. StarFOX overcomes these limitations by applying the angles-only Absolute and Relative Trajectory Measurement System (ARTMS), which integrates three novel algorithms. Image Processing detects and tracks multiple targets in images, using multi-hypothesis methods and domain-specific kinematic modeling, and computes target bearing angles. Batch Orbit Determination computes initial swarm orbit estimates from bearing angle batches, via iterative batch least squares and sampling of the weakly observable target range. Sequential Orbit Determination leverages an adaptive, efficient unscented Kalman filter with nonlinear models to refine swarm state estimates over time. Multi-observer measurements shared over an intersatellite link are seamlessly fused to enable robust absolute and relative orbit determination. StarFOX flight data and telemetry presents the first demonstrations of autonomous angles-only navigation for a satellite swarm, including multi-target and multi-observer relative navigation; autonomous initialization of navigation for unknown targets; and simultaneous absolute and relative orbit determination. Relative positioning errors of 0.5% of target range are achieved for a single observer under challenging measurement conditions, reduced to 0.1% with multiple observers. Results demonstrate promising performance with regards to ongoing StarFOX campaigns and the application of angles-only navigation to future distributed missions.

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 Situational Awareness (SSA)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 simi-

lar 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 on-board 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 research in angles-only navigation for spacecraft and several prior attempts to conduct angles-only navigation in flight. These include the Orbital Express mission,¹² the Advanced Rendezvous using GPS and Optical Navigation (ARGON) experiment,^{13,14} and the Autonomous Vision Approach Navigation and Target Identification (AVANTI) experiment.^{15,16} Of these, ARGON and AVANTI are well-documented in literature.

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

AVANTI (2016) was conducted during the FIRE-BIRD mission of the German Aerospace Center. It encompassed the rendezvous of the BIROS microsatellite and an ejected nanosatellite from 13 km to 50 m of Inter-Satellite Distance (ISD). In comparison, AVANTI was much more autonomous and image processing, relative navigation, and maneuver planning were all performed on board. Targets could be autonomously tracked with less reliance on orbit knowledge from two-line elements; an extended Kalman filter was used to perform ROE state estimation; and maneuver planning employed a model predictive control strategy for optimization. Passive safety was also enforced via eccentricity/inclination-vector separation. Notably, non-ideal conditions resulted in significant bearing angle measurement outages, as well as a degradation of attitude determination perfor-

mance at closer ranges when target brightness outshone background stars in the FOV.

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

The goal of the Starling Formation-flying Optical eXperiment (StarFOX), as described in this paper, is to remove these limitations entirely. StarFOX is one of four experimental payloads of the NASA Starling mission managed by the NASA Ames Research Center,¹⁷ and is visualized in Figure 1. Starling consists of four 6U CubeSats in LEO and aims to increase the readiness of four enabling technologies for autonomous spacecraft swarms: operations and decision-making, communications and networking, maneuver planning and execution, and absolute and relative navigation.



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

Specifically, StarFOX applies the angles-only Absolute and Relative Trajectory Measurement System (ARTMS). ARTMS is a self-contained architecture^{18,19} that provides distributed, autonomous, scalable navigation capabilities for DSS orbiting an arbitrary central body, without reliance on maneuvers or external measurement sources. Starling was launched in July 2023, and over the ensuing months, a series of StarFOX experiment blocks has explored 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 two different swarm formations. The experiment campaign consequently builds

upon ARGON and AVANTI for a more wide-ranging and ambitious demonstration of spaceborne angles-only navigation in orbit.

This paper provides an overview of StarFOX operations during the first experiment phase, as well as initial flight results obtained by 1) post-processing of flight imagery and sensor data via ARTMS flight code running on the ground, and 2) ARTMS operating on board the Starling swarm, with comparisons to pre-flight expectations. These are the first demonstrations of autonomous angles-only navigation for a satellite swarm, including multi-target and multi-observer relative navigation; autonomous initialization of navigation for unknown targets; and simultaneous absolute and relative orbit determination. Though initial results are promising, significant challenges were also encountered, motivating key lessons learned for future StarFOX experiments.

Following this introduction, the angles-only navigation problem is defined. Then, ARTMS and its algorithms are described in more detail, followed by an overview of the Starling mission and hardware. Next, the paper describes StarFOX experiment plans and operations. Subsequently, on-ground and on-orbit results for several key navigation scenarios are presented, followed by lessons learned and conclusions.

MATHEMATICAL BACKGROUND

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 $\hat{x}^{\mathcal{R}}$ (directed along the observer's absolute position vector); $\hat{z}^{\mathcal{R}}$ (directed along the observer's orbital angular momentum vector); and $\hat{y}^{\mathcal{R}} = \hat{z}^{\mathcal{R}} \times \hat{x}^{\mathcal{R}}$.²⁰ Similarly, define a frame \mathcal{W} using $\hat{y}^{\mathcal{W}}$ (directed along the observer's velocity vector); $\hat{z}^{\mathcal{W}} = \hat{z}^{\mathcal{R}}$; and $\hat{x}^{\mathcal{W}} = \hat{y}^{\mathcal{W}} \times \hat{z}^{\mathcal{W}}$. \mathcal{W} only differs from \mathcal{R} by a rotation of the observer flight path angle ϕ_f about $\hat{z}^{\mathcal{R}}$ with $\phi_f \approx 0$ in near-circular orbits.²⁰ Finally, define the observer VBS coordinate frame \mathcal{V} consisting of orthogonal basis vectors $\hat{x}^{\mathcal{V}}$, $\hat{y}^{\mathcal{V}}$, $\hat{z}^{\mathcal{V}}$ where $\hat{z}^{\mathcal{V}} = \hat{x}^{\mathcal{V}} \times \hat{y}^{\mathcal{V}}$ 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 $\hat{z}^{\mathcal{V}}$ is aligned with $\pm \hat{y}^{\mathcal{W}}$. Figure 2 illustrates this scenario.

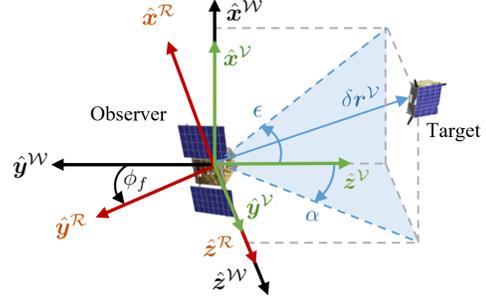


Figure 2: Definition of coordinate frames and bearing angles with VBS pointing in the $-\hat{y}^{\mathcal{V}}$ direction.

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 o to target t is¹⁸

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

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

Measurements and states are also referenced with respect to an inertial reference frame centered on an arbitrary central body, denoted \mathcal{I} . The inertial frame 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 $\delta \mathbf{r}^{\mathcal{I}} = {}^{\mathcal{V}}\mathbf{R}^{\mathcal{I}} \delta \mathbf{r}^{\mathcal{V}}$ where ${}^{\mathcal{V}}\mathbf{R}^{\mathcal{I}}$ 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.²¹ Other relevant rotations ${}^{\mathcal{R}}\mathbf{R}^{\mathcal{I}}$ and ${}^{\mathcal{W}}\mathbf{R}^{\mathcal{I}}$ can be computed using the observer's absolute orbit estimate. Note that in practice, VBS measurement availability is affected by optical visibility constraints such as eclipse periods.

System State

Because angles-only navigation is characterized by weak observability,⁹ proper selection of the state parameterization is crucial to maximize estimation accuracy and robustness to errors. To meet these objectives, ARTMS leverages quasi-nonsingular absolute Orbit Elements (OE) and Relative Orbit Elements (ROE).

The absolute orbit α is defined as

$$\alpha = [a \ e_x \ e_y \ i \ \Omega \ u]^T = [a \ e \cos \omega \ e \sin \omega \ i \ \Omega \ \omega + M]^T \quad (3)$$

where $a, e, i, \Omega, \omega, M$ are the Keplerian orbit elements. The relative orbit of each target detected by the on-board sensor, denoted $\delta \alpha$, is described by the quasi-nonsingular ROE adopted by D'Amico.²² Each of these ROE is function of the orbit elements of the target t and

observer o as given by

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

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

These state definitions have been previously studied in literature, resulting in development of accurate analytical dynamics models.^{23,24} 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.²⁵ 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.

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.^{18,26-28} However, the initial StarFOX experiment phase only explores orbit estimation to reduce computation costs and system complexity. Auxiliary state estimation will be implemented as part of a StarFOX extended mission.

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

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

$$\mathbf{d}^{\mathcal{R}} = [d_x^{\mathcal{R}} \quad d_y^{\mathcal{R}} \quad d_z^{\mathcal{R}}]^{\top} \quad (6)$$

where $G \in \mathbb{R}^{6 \times 3}$ is the GVE matrix²⁹ 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.³⁰ In flight, ARTMS typically applies a 10x10 GM01S² spherical harmonic gravity model, a Harris-Priester atmosphere model with cannonball drag ($C_D \approx 0.015$), and a 30 second RK4 step. Analytic dynamics models for the mean OE which include the effects of J_2 gravity are used within ARTMS when computational efficiency is paramount.

With regards to relative dynamics, the ROE provide useful geometric intuition regarding target relative motion. The curvilinear position vector of a target in the observer's RTN frame is defined as $\delta \mathbf{r}_{\text{curv}}^{\mathcal{R}} = (\delta r, a_o \Theta, a_o \Phi)$.³¹ 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 captures the effects of orbit curvature with improved accuracy and can be mapped back to rectilinear coordinates via

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

As first shown by D'Amico,²² there exists a linear mapping between the ROE and $\delta \mathbf{r}_{\text{curv}}^{\mathcal{R}}$, defined as

$$\delta \mathbf{r}^{\mathcal{R}} \approx a_o \begin{bmatrix} \delta a - \delta e_x \cos u_o \\ -1.5\delta a + \delta\lambda + 2\delta e_x \sin u_o - 2\delta e_y \cos u_o \\ \delta i_x \sin u_o - \delta i_y \cos u_o \end{bmatrix} \quad (8)$$

The mapping was later extended to eccentric orbits by defining eccentric ROE³² $\delta\alpha^* = [\delta a, \delta\lambda^*, \delta e_x^*, \delta e_y^*, \delta i_x, \delta i_y]^{\top}$, which revert to the quasi-nonsingular ROE for $e_o \approx 0$.

Figure 3 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.

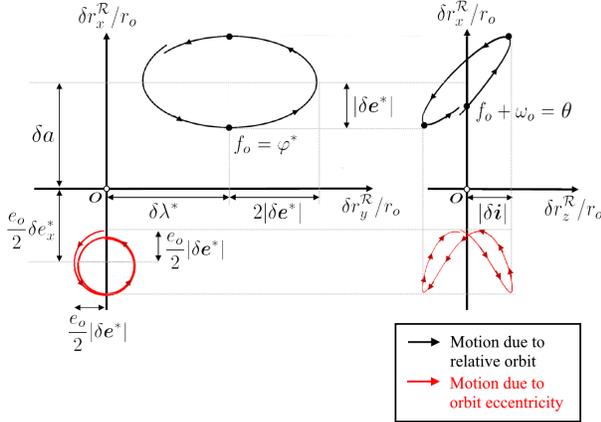


Figure 3: Components of target relative motion in the $\hat{x}^R-\hat{y}^R$ (RT) and $\hat{x}^R-\hat{z}^R$ (RN) planes.¹⁸

NAVIGATION ARCHITECTURE

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 4 presents a notional illustration of a four-spacecraft DSS using ARTMS.

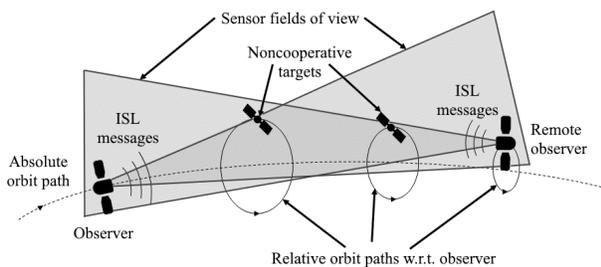


Figure 4: Notional illustration of ARTMS observers and targets for a four-spacecraft system (not to scale).

A high-level overview of ARTMS as implemented in StarFOX flight software is shown in Figure 5. The architecture of three core modules based on algorithms recently developed at Stanford’s Space Rendezvous Laboratory (SLAB): IMage Processing (IMP),³³ Batch Orbit Determination (BOD),¹¹ and Sequential Orbit Determination (SOD).¹⁸ 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, and receives ARTMS telemetry. If the spacecraft is equipped with a GNSS receiver, it optionally provides Position/Velocity/Time (PVT) navigation solutions to replace less timely orbit estimates from the ground. This constitutes the only a-priori information needed, i.e. to initialize, each observer must possess a coarse estimate of its own orbit at a single epoch.

The IMP module uses a coarse estimate of the observer’s orbit and VBS images to produce batches of bearing angles and corresponding uncertainties to all visible targets without 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.

Overall, ARTMS provides real-time orbit estimates for the host spacecraft and each target detected by the onboard sensor. The only hardware requirements posed on the spacecraft are that it must have a VBS and an ISL. 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.

Image Processing

The objective of the IMP module is to produce batches of time-tagged bearing angle measurements to each target, using images from the onboard VBS. This is accomplished in two phases. First, each incoming image is processed and reduced to a set of inertial bear-

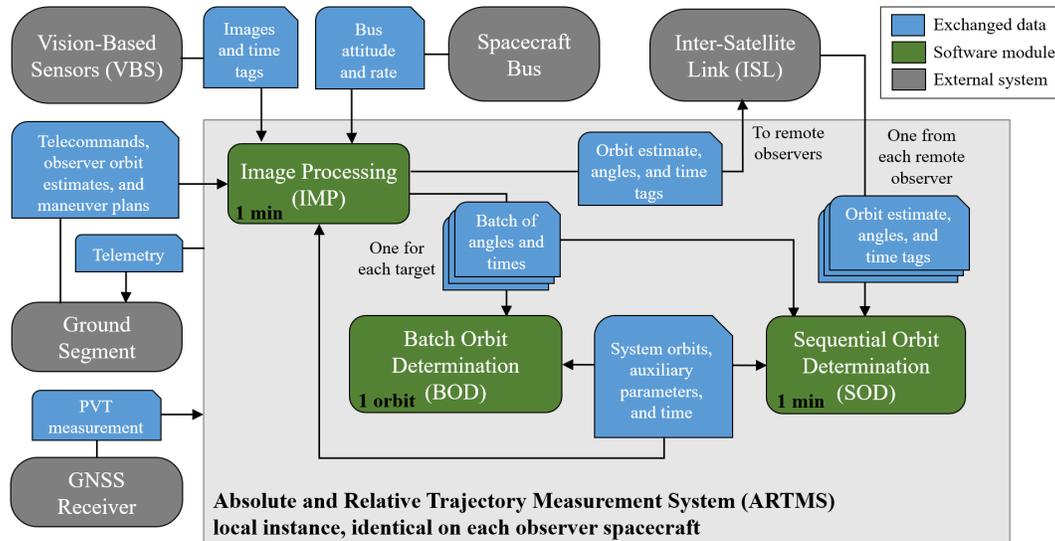


Figure 5: A high-level overview of the ARTMS flight architecture.

ing angles that may correspond to resident space objects. Second, these candidate bearing angles are used to track known targets and detect new targets using an approach inspired by Multi-Hypothesis Tracking (MHT).³⁴ Figure 6 presents an example image input, with typical sample rates of 60-120 seconds during StarFOX in LEO, equating to 50-100 images per orbit.

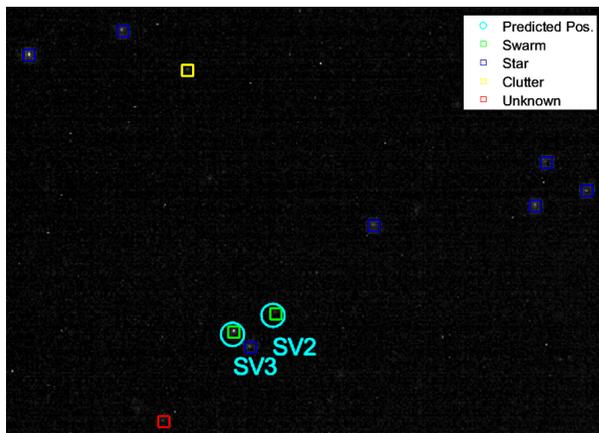


Figure 6: A Starling flight image with point sources detected and identified by IMP.

The first phase of IMP uses a set of existing algorithms. An iterative weighted center of gravity centroiding algorithm³⁵ is used to simplify the raw image into a list of pixel cluster centroids. Centroids are converted to unit vectors in the VBS frame using the calibrated sensor model. Next, the Pyramid star identification algorithm³⁶ is applied to remove Stellar Objects (SO) from the list of pointing vectors. Uncatalogued SO are detected by considering objects with unchanging inertial unit vec-

tors between images. Similarly, camera hotspots are removed by considering objects with unchanging pixel coordinates. The VBS attitude is computed from the pointing vectors to identified stars in the inertial and sensor frames using the q-method.³⁷ The remaining minimalistic set of inertial unit vectors (and equivalent bearing angles in \mathcal{V}) likely correspond to known targets or other unknown objects in the FOV.

In the second phase, these measurement candidates must be consistently assigned to new targets or targets that are currently being tracked without requiring a-priori relative orbit knowledge. To accomplish this, the IMP module employs the new Spacecraft Angles-only Multitarget tracking System (SAMUS),³³ which has only two key requirements: 1) a coarse estimate of the observer's absolute orbit must be provided, and 2) targets must not perform unknown translational maneuvers during the tracking period. SAMUS is valid for eccentric orbits and has been specifically 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.

SAMUS applies the core concept of MHT in that as measurements arrive, many simultaneous hypotheses are maintained as to how they can be associated into target tracks. The algorithm converges towards the correct hypothesis over time, to enhance robustness of its outputs. MHT is chosen as a basis because it is mature and demonstrably accurate, with its primary disadvantage being the need to frequently and heuristically prune hypotheses for real-time computation.³⁸ To overcome this, SAMUS applies domain-specific knowledge to develop

precise pruning criteria.

Recall Equation 7, which maps the ROE to a target’s curvilinear position vector in the observer’s RTN frame. In the map, u is the only quickly-varying term, while all other terms vary slowly in the presence of perturbations. Thus, target motion is described by periodic, parametric functions with known form. Even if the specific ROE in Equation 7 are unknown, its form provides expectations regarding target motion that can be leveraged.

Given a set of past bearing angles measurements in a track and corresponding estimates of the observer’s absolute orbit, Equation 7 can be rearranged into a pair of separable linear systems in azimuth and elevation,³³ as in

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

where $x_{1,\dots,6}$ are scaled ROE equivalents in bearing angle space.³³ 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.

To assess which hypotheses are physically reasonable, SAMUS applies a set of kinematic rules³³ derived from the parametric motion model. Only tracks which pass all rules are propagated. Briefly, the rules are summarised as: 1) track velocities must be below a set maximum, 2) track velocities must be consistent over time, 3) tracks should generally not feature acute angles, 4) tracks should turn in a consistent direction, 5) new measurements must be close to the predicted measurement. Their application greatly increases efficiency of MHT by preventing formation of unlikely tracks.

When multiple tracks pass all rules, SAMUS scores propagated tracks via ten criteria which assess how well each fulfills the expectations of Equation 7, the predicted measurement, and their prior motion. In contrast to traditional MHT methods – which often rely on a single Mahalanobis distance 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. By using a larger set of metrics, consensus supports the correct choice over time, even if some temporarily support incorrect hypotheses. Additionally, scoring does not require probabilistic estimates of false alarm densities or target decay rates, which are not easily obtainable for spacecraft.

To initialize new tracks, SAMUS employs the Density-Based Spatial Clustering of Applications with Noise (DBSCAN) algorithm.³⁹ DBSCAN clusters re-

quire $\geq n_D$ points within small radius ϵ_D . Because targets are in similar orbits to observers, their velocities compared to other objects in the FOV are low. Previously untracked targets are initialized by applying DBSCAN to the merged set of unidentified measurements from the past several images, and applying the SAMUS kinematic rules to found clusters.

Finally, tracking is often interrupted by orbit eclipse periods. To connect shorter tracks on either side of an eclipse, the aforementioned linear system fit is computed for every possible set of paired tracks. The combination of compatible pairs which produces the least fitting residuals is chosen as output.

SAMUS is also able to cooperate with SOD and apply target state knowledge (if available) for improved computation cost. Target state estimates from SOD are propagated into the current epoch to provide predicted track measurements. The kinematic rules are replaced by a validity region around the prediction, computed via an unscented transform of the target state covariance. The Mahalanobis distance between predicted and assigned measurement is employed for track scoring. Figure 7 presents an overview of core SAMUS operations.

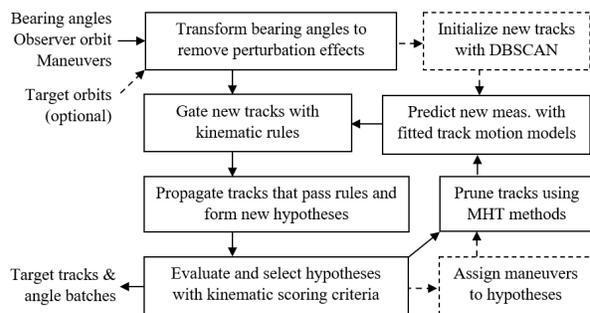


Figure 7: SAMUS algorithm summary and core sequence of operations.

Batch Orbit Determination

The BOD module must be able to produce orbit estimates for the local swarm with sufficient accuracy to initialize the SOD module using only a coarse estimate of the observer’s orbit and batches of bearing angles to each target from IMP. Nominally, the module operates on measurement batches of [50, 200] measurements across 1-2 orbits, per target. State estimation is accomplished using an algorithm that sequentially estimates the relative orbits of each target while simultaneously refining the estimate of the semimajor axis of the observer’s orbit.¹¹ When used in LEO, the algorithm applies a fully analytical dynamics model including the earth oblateness J_2 perturbation developed by the authors to minimize computation cost.²³ For each target, the estimation process is a four-step procedure.

First, a one-dimensional family of state estimates is computed for user-specified values of $\delta\lambda$. The set of possible $\delta\lambda$ values is selected based on system limitations such as sensitivity of the VBS. The estimated state vector \mathbf{x}_{est} corresponding to each $\delta\lambda$ value is initialized with all other ROE being zero. Then, iterative batch least squares refinement is performed for the ROE and semimajor axis.

Define the batch of true angle measurements from IMP as \mathbf{z}_{meas} for measurement epochs \mathbf{t} , and a batch of modeled angle measurements $\mathbf{z}_{\text{model}}$, computed by propagating the estimated target state to each epoch in \mathbf{t} . The difference between measured and modeled angles is

$$\Delta\mathbf{z} = \mathbf{z}_{\text{meas}} - \mathbf{z}_{\text{model}}(\mathbf{x}(t_{\text{est}}), \mathbf{t}) \quad (10)$$

where t_{est} is the BOD estimation epoch. Subsequently, the state estimate update

$$\Delta\mathbf{z} = \mathbf{S}_{\text{est}}(\mathbf{x}(t_{\text{est}}), \mathbf{t})\Delta\mathbf{x}_{\text{est}} \quad (11)$$

$$\mathbf{x}_{\text{est}} \leftarrow \mathbf{x}_{\text{est}} + \Delta\mathbf{x}_{\text{est}} \quad (12)$$

can be performed, where $\mathbf{S}_{\text{est}}(\mathbf{x}(t_{\text{est}}), \mathbf{t})$ is a sensitivity matrix containing the partial derivatives of measurements with respect to state components, computed via linearization of the bearing angle measurement model. 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.

Second, the final state estimate is selected as the candidate state estimate with the smallest measurement residual vector $\Delta\mathbf{z}$. Figure 8 shows the norms of the converged measurement residual vectors for each candidate value of $\delta\lambda$ in a single test case of the BOD algorithm. The line-sampling procedure for $\delta\lambda$ aims to maximize the observability of target range during initialization.

Third, the measurement noise matrix for each measurement (denoted $\mathbf{R}_{\text{sensor}}$) is estimated using the measurement residuals corresponding to the final state estimate. 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.

Fourth, the covariance for estimated state components \mathbf{P}_{est} is computed as given by

$$\mathbf{P}_{\text{est}} = \mathbf{Y}_{\text{est}}^* (\mathbf{N}\mathbf{R}_{\text{sensor}} + \mathbf{Y}_{\text{prior}}\mathbf{P}_{\text{prior}}\mathbf{Y}_{\text{prior}}^T) \mathbf{Y}_{\text{est}}^{*T} \quad (13)$$

where $\mathbf{Y}_{\text{est}}^*$ is the pseudoinverse of the measurement sensitivity matrix for estimated state components, $\mathbf{Y}_{\text{prior}}$ is the measurement sensitivity matrix for a-priori information (orbit elements other than the semimajor axis, ballistic coefficients, sensor biases, etc.), $\mathbf{P}_{\text{prior}}$ is the uncertainty of the a-priori information, and N is the number

of provided bearing angle measurements. This formulation allows the BOD module to seamlessly transition between domains where uncertainty is driven by sensor performance and by errors in the a-priori information. The ROE estimates for each target are then appended to the refined estimate of the observer's absolute orbit, forming a complete estimate of the state of the local swarm.

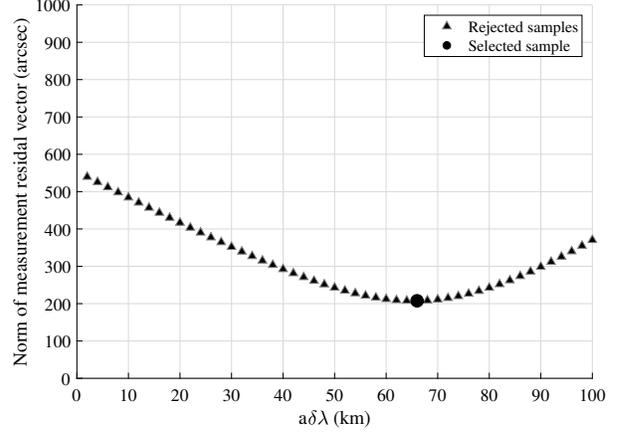


Figure 8: Behavior of converged measurement residuals for different $\delta\lambda$ samples in the BOD module.

Sequential Orbit Determination

The SOD module continually refines estimates of the orbits of all spacecraft in the local swarm as well as auxiliary parameters (e.g. differential clock offsets) by seamlessly fusing measurements from all observers transmitted over the ISL. SOD is based on a new adaptive, efficient Unscented Kalman Filter (UKF) with fully nonlinear dynamics and measurement models. In comparison to an extended Kalman filter, the UKF preserves higher order moments in the probability distribution, which enables maneuver-free convergence using angles-only measurements from a single observer, without excessive computational cost.¹⁸

Three additional features are included to maximize performance using measurements from multiple observers. 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.⁴⁰ Third, measurements from remote observers are assigned to tracked targets using selection criteria based on the Mahalanobis distances between the estimated bearing angles to each target and each candidate measurement. Let σ_{jk} denote the Mahalanobis distance between the j th measurement from the remote observer and the predicted

measurement to the k th target tracked by the local observer (which accounts for all relevant state uncertainties). To minimize erroneous assignments, the j th measurement from the remote observer is assigned to the k th target if three conditions are satisfied:

- 1) $\sigma_{jk} \leq \epsilon_{assign}$
- 2) $\sigma_{lk} \geq \epsilon_{ambig} \forall l \neq j$
- 3) $\sigma_{jp} \geq \epsilon_{ambig} \forall p \neq k$

where ϵ_{assign} and ϵ_{ambig} are user-specified parameters that satisfy $\epsilon_{ambig} > \epsilon_{assign} > 0$. These conditions ensure 1) the measurement is close to the modeled measurement of the target using the current state estimate, 2) there is no other candidate measurement that fits the estimated state of the target, and 3) there is no other target with a state estimate that fits the measurement.

Figure 9 includes conceptual illustrations of four possible cases of modeled and observed measurements from a remote observer which (from left to right) show all conditions satisfied and violations of Condition 1, Condition 2, and Condition 3, respectively. Together, these conditions ensure that measurements are only assigned when the observed and modeled measurements uniquely agree with a statistical certainty determined by the values of ϵ_{assign} and ϵ_{ambig} . The values of these parameters should be selected based on the expected number of targets, relative motion geometry, sensor noise, and available orbit knowledge for general swarming missions. However, for scenarios similar to StarFOX, the authors have found that setting $\epsilon_{assign} = 3$ and $\epsilon_{ambig} = 6$ provide robust measurement assignment performance.

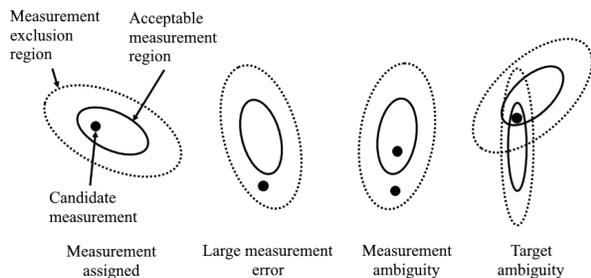


Figure 9: Conditions in which all measurement assignment criteria are satisfied (leftmost) and conditions that violate each criteria (right).

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: 1) the Mahalanobis distance between the most recent BOD relative state estimate and current SOD relative state estimate is above a threshold, 2) the number of successive filter update steps which did not feature a measurement update is above a threshold, 3) the number of successive filter measurement update steps which

featured a large pre-fit measurement residual is above a threshold, 4) the SOD state estimate is outside a valid range (as specified by the user). To re-initialize, the SOD module may either use ground-provided orbit estimates or the most recent BOD state estimate.

ON-ORBIT TESTBED

StarFOX is a core payload of the NASA Starling swarm mission, which consists of four 6U CubeSats. Starling was initially proposed in 2018¹⁷ as a testbed for autonomous swarming technologies, with a focus on four separate experiment payloads: autonomous networking, autonomous decision-making, autonomous maneuvering, and autonomous navigation. Development of the StarFOX navigation experiment began in 2019, and integration and testing of the ARTMS flight software continued until the end of 2022.¹⁹ The swarm was launched into an approximately sun-synchronous low Earth orbit on July 17, 2023. The primary mission phase was expected to allow for approximately six months of experiments, with the spacecraft transitioning through multiple formations. Nominal experiment operations commenced in November 2023 and continued until May 2024.

Bus Hardware

For StarFOX, each CubeSat carries two Blue Canyon Technologies Nano Star Trackers⁸ (NST) aligned in antiparallel directions, either of which may collect imagery. Parameters of the NST are provided in Table 1. During each experiment, one NST is chosen as the experiment camera and is pointed in the velocity or anti-velocity direction, to image swarm members. Images were provided to ARTMS with a cadence of 60-120 seconds. The other NST provides a backup spacecraft attitude solution.

Table 1: Intrinsic parameters of the NST.

Parameter	Value
Image Size (pixels)	1280 × 1024
FOV (°)	12 × 10
Pixel Size (μm)	5.3
Focal Length (mm)	30
Pixel Space	12-bit (grayscale)

StarFOX also receives inputs from the onboard GPS receiver, which optionally provides PVT solutions, and from the ISL, through CesiumAstro S-band software-defined radios. If the ISL is active, each spacecraft transmits a message to all other swarm members every 60 seconds and listens for messages from all other members.

The payload processor for StarFOX is a Xiphos Q7S running at approximately 700 MHz. This payload processor is separate from the primary flight processor and

similarly, the ARTMS flight software is kept separate from the higher-level core flight software, to avoid impacting nominal mission operations. Finally, each satellite also possesses a cold gas propulsion system with four separate nozzles, used to conduct swarm phasing and re-configuration maneuvers.

Figure 10 shows the four Starling spacecraft during integration at the NASA Ames Research Center. The spacecraft are denoted SV1, SV2, SV3 and SV4 or more affectionately as Inky, Blinky, Pinky, and Clyde.



Figure 10: The four Starling CubeSats during integration. Credit: NASA/Dominic Hart (2022).

Swarm Geometry

Swarm geometry for the Starling mission was determined by the differing needs of flight experiments and in-flight safety considerations. For StarFOX, the following requirements were specified:

1. The ISD between swarm members shall not exceed crosslink communication range.
2. The ISD shall not exceed the limit for detection of targets in star tracker images.
3. ISD shall remain greater than a minimum safe distance.
4. Instantaneous bearing angle measurements of targets shall not remain constant (i.e. there should be observable relative motion).
5. Relative motion between spacecraft shall include unique variations in the radial/normal directions.
6. Relative motion ellipses shall lie entirely within the VBS FOV when the VBS boresight is aligned with the (anti)-velocity direction.

In order to provide varying swarm geometry for StarFOX, two distinct formations were used: an in-train (IT) formation and a passive safety ellipse (PSE) formation. The swarm was initially placed into the IT formation, then transitioned into the PSE formation in March 2024.

Table 2 presents OE for SV4, and ROE for SV2 and SV3 (with respect to SV4), following the conventions of Equations 3 and 4. SV1 was not involved in primary

StarFOX operations. Figure 11 displays the relative motion geometry of the swarm in the radial-tangential and radial-normal planes of SV4, in curvilinear coordinates. The magnitude of relative motion is larger in the PSE formation to provide increased angles-only observability and passive safety, while still allowing target observation with a constant attitude profile.

Table 2: Orbit elements of the swarm.

ID	In-Train Formation (on 02/05/24 at 00:00:00 UTC)
SV4	$\alpha = [6945 \text{ km}, 0.0007, 0.0014, 99.4^\circ, -152.3^\circ, -41.8^\circ]$
SV2	$\delta\alpha = [21, -124350, 110, 202, 79, 1005] \text{ m}$
SV3	$\delta\alpha = [-1, -79328, 42, 452, 36, 827] \text{ m}$
ID	Passive Safety Ellipse (on 05/08/24 at 00:00:00 UTC)
SV4	$\alpha = [6943 \text{ km}, 0.0001, 0.0016, 99.4^\circ, -38.0^\circ, -153.0^\circ]$
SV2	$\delta\alpha = [23, -52266, -471, -523, 35, 1811] \text{ m}$
SV3	$\delta\alpha = [-30, 155320, -297, -315, 16, 1144] \text{ m}$

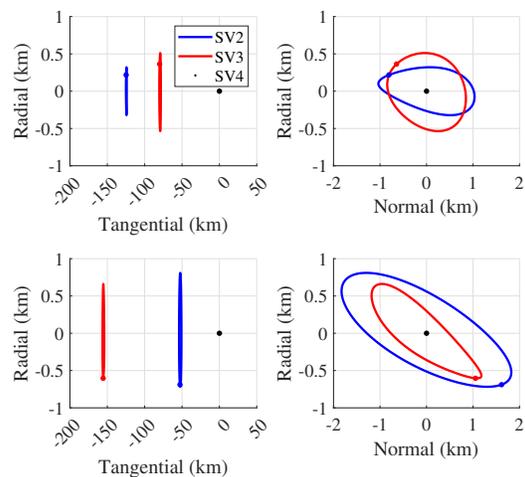


Figure 11: Swarm relative motion in the IT formation (upper) and PSE formation (lower).

FLIGHT OPERATIONS

The initial Starling mission was divided into six phases: 1) launch and deployment, 2) commissioning, 3) formation acquisition, 4) in-train formation, 5) passive safety ellipse formation, 6) extension. This paper concerns StarFOX experiment operations during Phases 4) and 5). Ongoing communication and propulsion difficulties affecting SV1 meant that it could not be reliably included in StarFOX experiments, and furthermore, its relative orbit was not consistently visible from the other swarm members. Experiments were therefore planned using SV2, SV3 and SV4 only.

Experiment Planning

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

1. Single observer, partially autonomous: initial target relative orbit estimates are provided by the ground and refined on board by IMP + SOD. GPS is used for observer absolute orbit estimation.
2. Single observer, fully autonomous: initial target relative orbit estimates are generated on board by IMP + BOD and refined on board by IMP + SOD. GPS is used for observer absolute orbit estimation.
3. Multi-observer, partially autonomous: multiple observers share measurements over the ISL to estimate target relative orbits with improved accuracy. GPS is used for observer absolute orbit estimation.
4. Multi-observer, fully autonomous: multiple observers share measurements over the ISL to estimate target relative orbits with improved accuracy. Inter-satellite bearing angles are used for observer absolute orbit estimation.

With reference to potential applications, Modes 1 and 3 are most applicable to cooperative swarm scenarios in Earth orbit; Mode 2 is most useful for SSA in which unknown targets must be detected and tracked; and Mode 4 is most useful for deep space applications in which GPS is unavailable.

All four navigation modes were investigated in flight. Fourteen experiment blocks were originally planned,¹⁹ differentiated by swarm geometry, visible targets, presence of maneuvers, presence of the crosslink, presence of GPS measurements, and presence of a-priori relative orbit knowledge. In total, eighteen StarFOX experiment blocks were executed, over a total of 34 experiment days. Experiment complexity increased throughout the schedule, with early experiments focusing on single-observer ground-assisted navigation and later experiments focusing on multi-observer autonomous navigation. Each swarm member could also be commanded independently, to allow multiple simultaneous experiments setups to be executed.

Experiment Execution

The typical sequence for a StarFOX experiment is as follows. Initial preparation involves the fusion of data from three sources: NASA Flight Dynamics (FD), NASA mission operators, and Stanford SLAB. FD provides predicted swarm ephemerides and planned maneuvers, which are used to provide each spacecraft with its own absolute orbit estimate for initializing navigation, and optionally, a-priori relative orbit estimates for its targets. SLAB provides a telecommand (TC) table for ARTMS which governs its behavior, e.g. activation of

individual ARTMS algorithms, multi-observer measurement use, and granular tuning of IMP, BOD and SOD. A NASA operator creates command sequences for the spacecraft bus to govern its with respect to StarFOX, including the attitude profile, imaging cadence, crosslink radio activation, ground station contact activities, and so on. These commands are scheduled in an Absolute Time Sequence plan, which is combined with the TC table into a final product. Subsequently, NASA's Mission Operations Systems (MOS) and Mission Operations Center perform command validation, using a hardware-in-the-loop simulation environment. Validated products are uplinked to the swarm.

Payload operations on board the spacecraft consist of several tasks. First, the spacecraft slews to a pre-defined attitude to ensure consistent swarm visibility, i.e. by aligning its star tracker with the (anti)-velocity direction. The TC table is unpacked and verified and regular experiment operations commence, including taking and copying star tracker images, running ARTMS applications, and saving ARTMS outputs as telemetry (TM). Nominal sample times for StarFOX are 60 seconds for IMP, 90 minutes for BOD, and 60 seconds for SOD. Note that StarFOX operations may be interrupted by various tasks onboard, such as swarm maneuvers and ground passes (which affect satellite attitude) or Globalstar signalling (which reduces GPS accuracy). Experiments continue for ~24 hours until a daily payload reset.

On-board TM is downlinked during regular ground passes and is provided to NASA Ground Data Systems. Raw TM from the StarFOX payload and satellite bus is made immediately available to MOS, and is processed to generate reports for payload experiment teams.

FLIGHT RESULTS

Flight results for StarFOX are generated in two ways. In the first case, ARTMS produces telemetry while operating in orbit, and telemetry is downlinked and visualized on the ground. In the second case, ARTMS inputs (i.e., star tracker images and metadata, GPS data, timing data, and ISL data) are downlinked, and are processed in a flight-like manner using ARTMS flight software running on the ground. The second case is a "digital twin" of the swarm which aims to operate as similarly as possible to ARTMS in flight. The digital twin was also used for diagnosing and troubleshooting ARTMS during the mission.

Navigation performance is assessed by comparison to "ground truth" orbit states, produced by NASA FD via batch processing of GPS flight data. Expected position errors and uncertainties in the ground truth are significantly less than 10 meters (1σ). High-level success criteria for StarFOX are summarized in Table 3. Subsequent

sections discuss flight results from four experiments in detail, corresponding to the four navigation modes.

Table 3: StarFOX mission success status.

Mission Requirement	Starling shall conduct inter-satellite relative position estimation.
Min. Success Criteria	Conduct ground-based relative position estimation using onboard measurements.
Full Success Criteria	Conduct onboard relative position estimation using onboard measurements.
Completion Status	Full success: onboard images used to produce onboard state estimates for multiple targets.

Relative Navigation: Single-Observer Multi-Target

Figure 12 presents in-flight estimation performance from SV4 telemetry on 05/16/2024. In this scenario, SV4 points its star tracker in the anti-velocity direction to image SV2 (at ~ 50 km range) and SV3 (at ~ 130 km range). Relative orbit estimates for SV2 and SV3 are initialized using stale information from the ground (approximately 3-4 days old). GPS is used to maintain SV4's absolute orbit estimate. Relative orbit estimation is performed using bearing angles from SV4 only. Images were taken every 60 seconds.

It can be seen in Figure 12 that SV4 successfully tracks and performs orbit estimation for multiple targets simultaneously, which is the first time this capability has been demonstrated in flight. It is also able to maintain a converged relative orbit estimate over a 21-hour period without requiring target maneuvers, which is another advancement compared to prior flight missions.

Performance displays expected trends for angles-only navigation. Error is primarily captured by the $\delta\lambda$ ROE, which corresponds to the weakly-observable target range. Error is also somewhat proportional to range, because at longer distances, the same sensor noise floor corresponds to a larger geometric uncertainty. Uncertainties (1σ) for $a\delta\lambda$ at the end of the experiment period are 1.3% for SV2 and 2.6% for SV3, as a percentage of target range. Uncertainties in other ROE are less than 0.02% of target range because these state components are more strongly observable. Estimation errors generally remain within 1σ bounds, indicating good health of the SOD filter. In this experiment, the filter assumes bearing angle measurement noise of $40''$ (1σ); a suggested smaller value of $20''^{8,19}$ was found to be insufficiently robust.

Results are also affected by unexpectedly poor target visibility, most obviously for SV3. In Figure 12, black dots indicate a detection of SV3 by SV4's star tracker, and fewer than 20 SV3 measurements were obtained over 12 orbits (whereas measurements SV2 were much more consistent). This is because the visual magnitude of the swarm was significantly dimmer than expected in orbit. Furthermore, the image signal-to-noise ratio was lower

than expected (e.g., due to the presence of hot pixels). As a result, targets at distances of more than 100 km were frequently lost in the image and could not be detected by IMP. Figure 13 illustrates this phenomenon: SV3 is consistently visible above background noise, whereas the more distant SV2 is not. Nevertheless, ARTMS was still able to leverage sparse measurements to maintain a reasonable orbit estimate for distant targets. This is evidence of the overall robustness of the architecture and its ability to maximize angles-only observability.

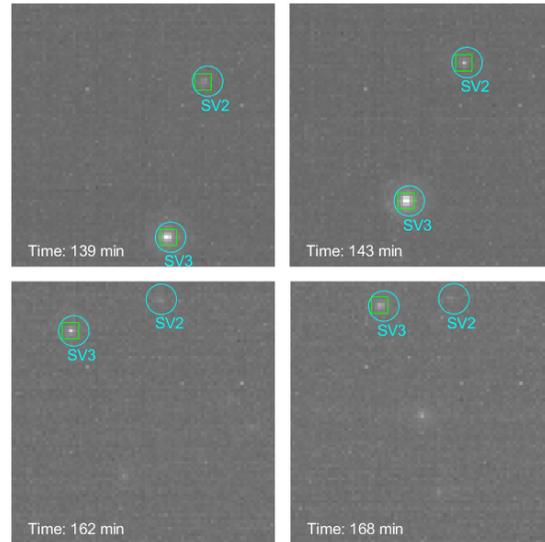


Figure 13: Target visibility in flight images from SV4. Green boxes indicate successful detection by IMP.

Relative Navigation: Multi-Observer Multi-Target

Figure 14 presents in-flight estimation performance from SV4 telemetry on 02/05/2024. In this scenario, SV4 points its star tracker in the anti-velocity direction to image SV2 (at ~ 130 km range) and SV3 (at ~ 80 km range). Relative orbit estimates for SV2 and SV3 are initialized using stale information from the ground (approximately 3-4 days old). GPS is used to maintain SV4's absolute orbit estimate. Relative orbit estimation is performed using bearing angles from multiple observers over the ISL, such that SV4 additionally receives crosslink measurements from SV2 (measuring SV3 and SV4) and SV3 (measuring SV2). Images were taken every 120 seconds.

The colored dots in Figure 14 mark instances when a measurement broadcast by a remote observer was used to update local state estimates on SV4. It can be seen that SV4 utilizes angles-only measurements from all three swarm spacecraft, which is the first time this capability has been demonstrated in flight.

The results also reveal several advantages of distributed multi-observer navigation. Firstly, SV4 only ob-

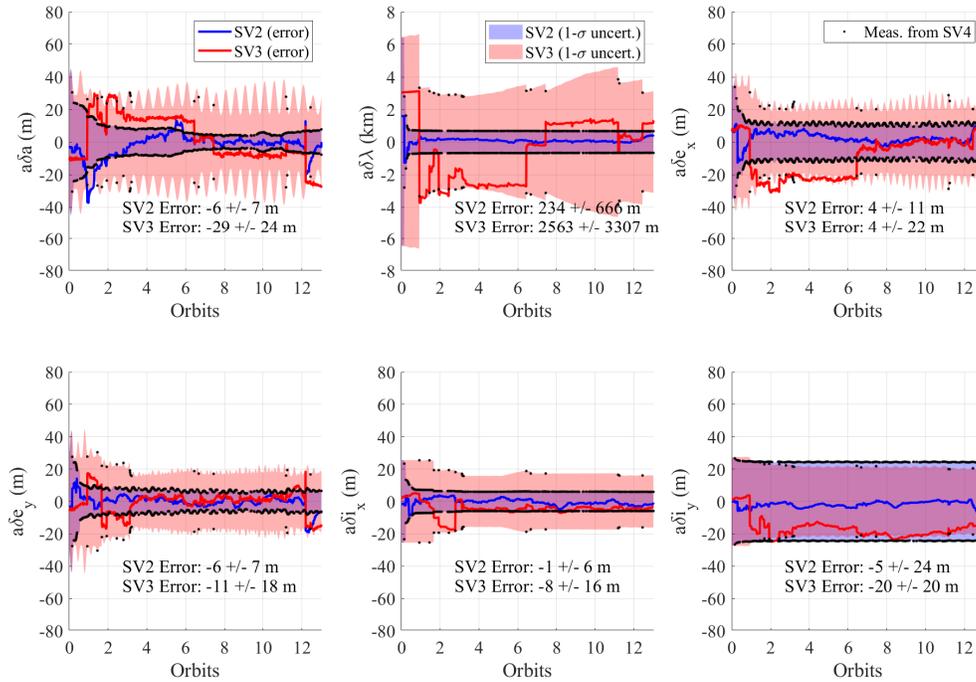


Figure 12: ROE state estimation errors and uncertainties (1σ) for the single-observer multi-target case.

tains 5 local measurements of SV2 throughout the entire period. It must instead rely on measurements of SV2 from SV3 (and/or SV3 from SV2) to update its on-board orbit estimate for SV2. Thus, multiple observers are able to provide added robustness when measurements are sparse. Secondly, state uncertainty is reduced compared to an equivalent single observer case. Uncertainties (1σ) for $a\delta\lambda$ at the end of the experiment period are 0.6% for SV2 and 0.6% for SV3, as a percentage of target range. Uncertainties in other ROE are less than 0.01% of target range. The additional geometric information provided by multi-observer measurements significantly improves angles-only observability.

Several anomalies are also evident in Figure 14, such as sudden decreases in the $a\delta\lambda$ uncertainty at several time instants (e.g. after approximately 6 orbits). This is due to unintended behavior of ARTMS when the SOD filter resets. As part of the ISL message, swarm members broadcast their absolute orbit estimate, which allows the receiving spacecraft to match any remote measurements to local state estimates. Typically, this absolute orbit estimate is taken from the on-board SOD filter estimate. However, when SOD re-initializes (e.g. due to a large number of skipped measurement updates), the SOD estimate becomes temporarily invalid, and the on-board GNSS solution is sent instead as backup. When SV4 re-

ceived this GNSS solution, it interpreted it as a measurement that could be used to update the relevant target's state estimate. GNSS provides improved observability in this scenario which led to a sudden improvement in that target's state uncertainty. (Remote GNSS navigation updates can be disabled within ARTMS but were left active during the 02/05/24 experiment.)

Earlier in the experiment, rapid covariance growth is observed after these instances, as relative orbit uncertainties stabilize after a GNSS update. Towards the end of the experiment, however, the steady-state uncertainty provided by multi-observer bearing angles is more consistently maintained. This implies the noise floor of multi-observer navigation is indeed at a similar level to the figure, even if the irregular GNSS updates may have sped up convergence.

Finally, note that there are certain periods when state estimation error ventures outside 1σ bounds, especially for $a\delta\lambda$ and $a\delta i_y$. However, errors do remain within 3σ bounds. This phenomena is believed to be caused by drifts in measurement time-tags between swarm members, discussed further in the following section.

Relative Navigation: Autonomous Initialization

Figure 15 presents in-flight estimation performance from SV2 telemetry on 05/08/2024. In this scenario, SV2

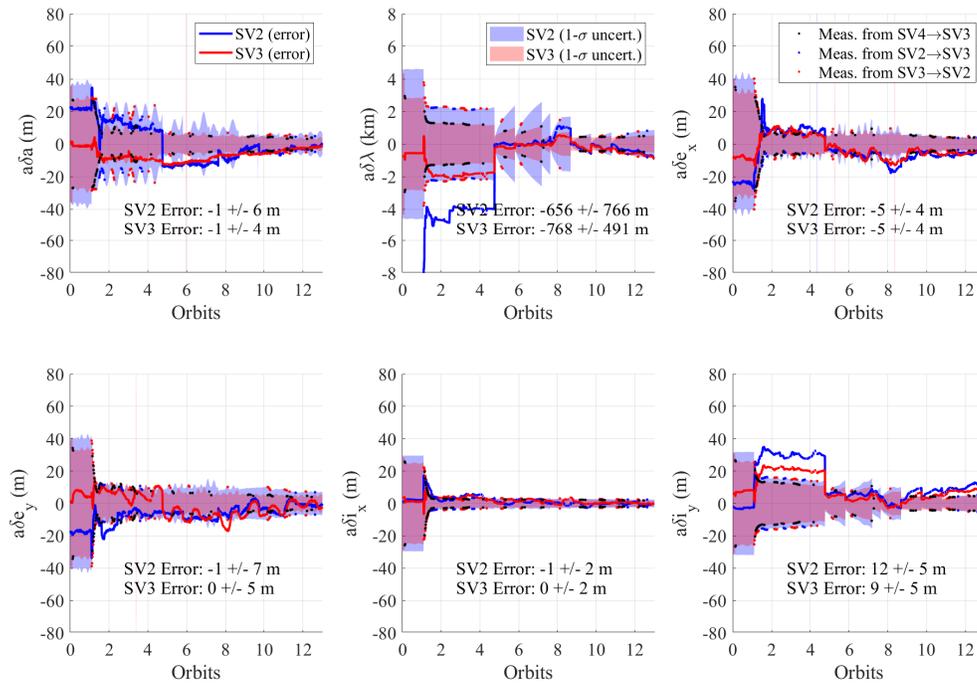


Figure 14: ROE state estimation errors and uncertainties (1σ) for the multi-observer multi-target case.

points its star tracker in the velocity direction to image SV4 (at ~50 km range). No a-priori relative orbit information for SV4 is provided. GPS is used to maintain SV4's absolute orbit estimate. SV2 begins by performing multi-hypothesis tracking with IMP to detect potential targets in the FOV. Every 1.5 hours, BOD attempts to compute a relative orbit initialization using IMP measurement batches. If the computed orbit is reasonable (i.e., is within state and uncertainty tolerances), it is used to initialize SOD and formal navigation commences. Images were taken every 60 seconds.

Figure 16 presents overlaid regions of interest from one orbit of images, as recorded in IMP telemetry. The initial ROE estimate was produced using a batch of 36 measurements, corresponding to the upper half of the ellipse in the figure. Figure 17 presents the measurement residuals for $\delta\lambda$ state samples and the output $\delta\lambda$ estimate, compared to the true ISD. Although the initial BOD range estimate possesses errors of 8 km (approximately 15% of the true range), this initial error is well-captured by the onboard uncertainty estimate. The convex form of the residual curve in Figure 17 also matches pre-flight expectations, as per Figure 8. Note that some difficulties were created by inconsistent image availability; of the 415 images expected to be available to IMP during this period, only 228 were successfully provided

to ARTMS by the satellite bus. Despite this, IMP was able to track the unknown target amid the background starfield and other transient RSO, and BOD was able to produce a reasonable initialization. This is the first time this capability has been demonstrated in flight.

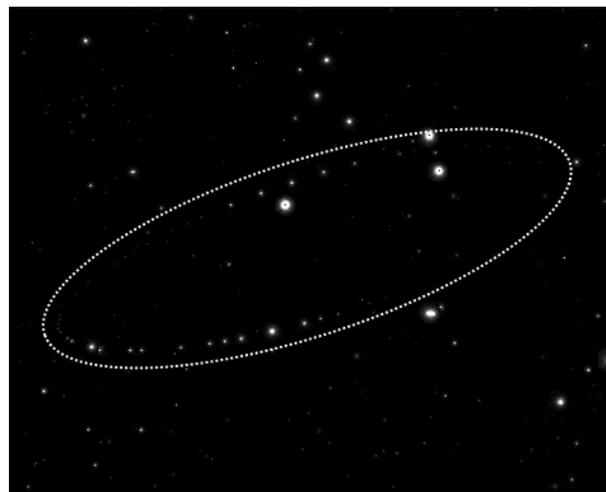


Figure 16: Overlaid regions of interest from one orbit of images, for SV2 viewing SV4 in the PSE formation. The point sources corresponding to the SV4's relative orbit are bordered by the ellipse.

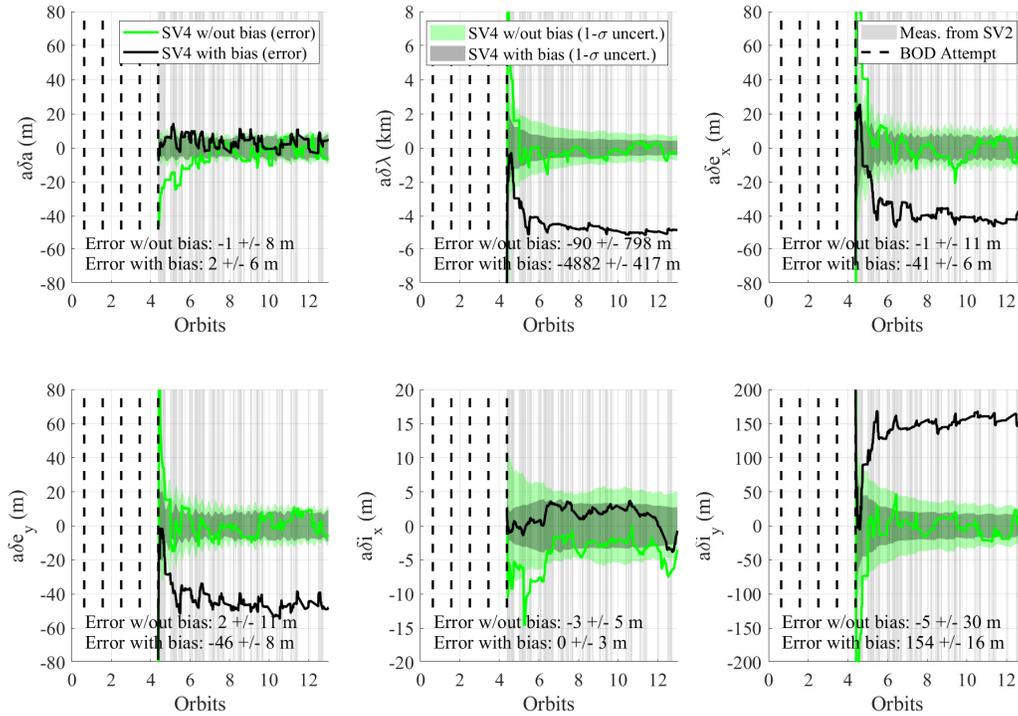


Figure 15: ROE state estimation errors and uncertainties (1σ) for the autonomous initialization case.

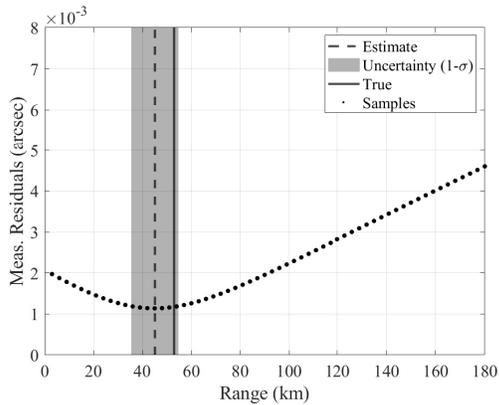


Figure 17: Initial target range estimation errors and uncertainties from BOD.

Figure 15 presents subsequent estimation performance. IMP + BOD autonomously initializes navigation on the fifth attempt, after approximately 4 orbits. The black line corresponds to estimation errors from in-flight telemetry, where it is clear that even if state uncertainties are converging, a strong bias is present in the state error for $\delta\lambda$, δe_x , δe_y , and δi_y . It was determined that this was produced by timing errors on board the satellite. The timestamps associated with star tracker images and meta-

data appeared to drift over time, and became unsynchronized with the primary clock (which was synchronized to GPS time). The green line in Figure 15 corresponds to estimation performance after the image timing errors are corrected, which removes the bias in all estimated ROE. The applied time correction for this experiment was -0.4 seconds, and was found to vary between spacecraft.

Absolute Navigation: Multi-Observer Multi-Target

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

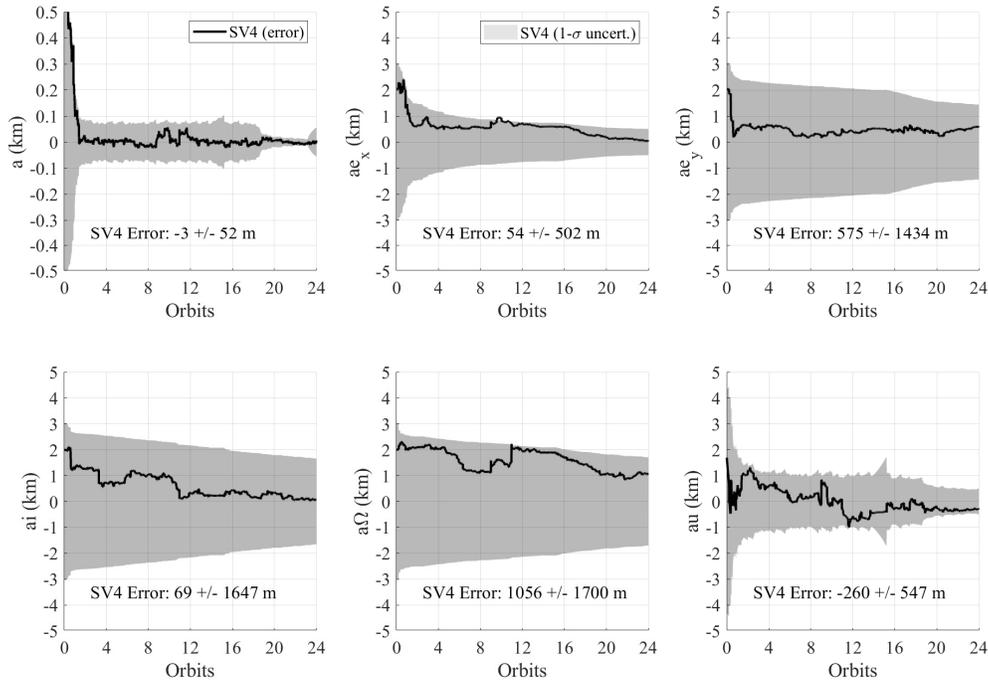


Figure 18: OE state estimation errors and uncertainties (1σ) for the multi-observer case.

Figure 19 presents absolute orbit estimation performance for SV4. Orbit element errors are normalized by the semi-major axis for geometric interpretation. Absolute orbit estimation using only inter-satellite bearing angles is particularly challenging due to the weak observability of the problem. However, if multiple observers are present and are obtaining measurements of a common target (and each other), enough geometric information is present to facilitate long-term absolute orbit convergence without external updates.^{19,28,41}

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

Figure 19 displays the ISL measurements received by SV4. The most consistently visible link was SV4 observing SV2, whereas other observations were more

sparse (due to the visibility issues previously described). The variation in measurement rates between seemingly equivalent spacecraft pairs (e.g. SV4 \rightarrow SV3 and SV3 \rightarrow SV4) illustrates an additional difficulty of applying optical measurements: visibility varies with specific orbit geometry and attitude, and image properties vary with specific VBS hardware. Scenario-specific tuning of image processing algorithms is therefore beneficial.

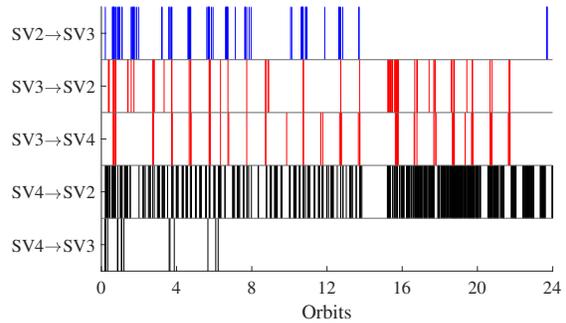


Figure 19: Measurements received by SV4. A bar indicates a measurement was received at that epoch.

Performance Summary

A series of experiment goals was defined for StarFOX prior to launch, informed by overall Starling objec-

tives and performance during pre-flight simulations and testing.¹⁹ These goals, and StarFOX outcomes during the initial experiment phase, are presented in Table 4. All thresholds were successfully achieved and the majority of goals were successfully achieved, through either in-flight experiments or on-ground flight data processing.

Objectives 6 and 7, which relate to swarm maneuvers, have not yet been discussed. Station-keeping maneuvers were executed during StarFOX experiments on 05/09/24 and 05/16/24, with the objective of observing a corresponding decrease in state uncertainty. Known maneuvers are expected to improve angles-only observability by aiding the filter in disambiguating target range.^{9,13} Although maneuvers were successfully tracked on board, the executed delta-vs were small (~ 0.01 m/s) and did not significantly improve state uncertainties as a result.

Lessons Learned

The flight experiences of StarFOX give rise to a number of lessons learned, which may be applied to any distributed mission employing vision-based angles-only navigation. The most critical of these is the overwhelming importance of realistic pre-flight simulation and validation, in software and hardware.

During validation of spacecraft flight software, it is typical to conduct a series of simulations which attempt to recreate the orbital environment as closely as possible. For StarFOX, the most crucial aspect is generation of realistic space images; significant attention was paid to producing appropriate imagery via an Optical Stimulator testbed with hardware-in-the-loop elements.^{19,42} Efforts were made to incorporate non-ideal measurement properties in regards to image noise, camera calibration errors, and measurement sparsity. Even so, in-flight conditions proved significantly more challenging than any assumptions made before launch. Target spacecraft were much dimmer than expected, sun intrusion into the FOV was not effectively modeled, and the image signal-to-noise ratio was too optimistic. In response, in-flight software updates were needed to enhance the adaptability of ARTMS to different optical conditions, via improved noise filtering and image normalization. This greatly improved performance in the second half of the mission. More realistic simulations may also have justified higher-level changes to the experiment design. For example, a more thorough determination of target visibility could have supported selection of a smaller nominal ISD, with strong benefits to visibility.

With respect to validation on spacecraft hardware, attention should be devoted to providing valid sensor inputs during the integration and testing (I&T) process. Unfortunately, resource constraints meant that image measurements could not be provided to ARTMS during

integration with the satellite bus, and thus, the majority of execution paths were not exercised. As a result, ARTMS encountered frequent software crashes during its first several experiments, caused by surpassing on-board memory limits. Though memory profiling had been performed pre-launch, the lack of realistic images during profiling meant that results were not representative. Second, as discussed, there was a variable drift between star tracker time-tags and the on-board clock aboard all swarm spacecraft. This could potentially have been discovered earlier if simultaneous image and GPS data was available during I&T. Instead, an in-flight software update was needed to allow user-specified time offsets for incoming star tracker data. More generally, this displays the importance of accurate time synchronization for navigation (especially if navigation is to be performed autonomously, e.g. without regular GPS time updates). The resulting state errors were significant and sources of time for different measurement types must be well-understood.

A final note is the importance of flexibility and robustness to overall system design. Efforts were made during the development of ARTMS to allow the system to adapt to differing conditions, in either an autonomous fashion (via different software execution paths, contingencies, and checks) or in a user-defined fashion (via tunable telecommand parameters). These decisions are what have enabled successful navigation even when in-flight conditions have proven particularly challenging.

Future Work

The Starling primary mission concluded in May 2024. However, StarFOX efforts will continue as part of Starling 1.5, a mission extension implemented via ground and onboard software updates. The capabilities of ARTMS will be augmented to include: 1) improved detection and tracking of unknown resident space objects for SSA, 2) faster orbit initialization for quickly evolving scenarios such as space rendezvous, 3) online estimation of auxiliary state components such as clock errors, drag coefficients and solar radiation pressure coefficients, 4) data fusion with inter-satellite crosslink ranging to achieve higher estimation accuracy, and 5) a fly-by inspection of a target satellite with control in the loop.

CONCLUSION

This paper presents an overview of the Starling Formation-flying Optical eXperiment (StarFOX), which is the first in-flight demonstration of autonomous angles-only navigation for a spacecraft swarm. StarFOX is a core experiment payload of the NASA Starling mission, managed by the NASA Ames Research Center. Starling

Table 4: Summary of StarFOX performance goals and initial flight results.

#	Objective	Threshold	Goal	Initial Result
1	Tracking multiple targets at one time	1 target	3 targets	2 targets (in flight)
2	Accuracy of relative position knowledge with one observer	Estimate produced	1% error relative to range	0.5% error relative to range (in flight)
3	Convergence time for relative position knowledge with one observer	Convergence observed	6 orbits	2 orbits (in flight)
4	Accuracy of relative position knowledge with multiple observers	Estimate produced	0.1% error relative to range	0.5% (in flight); 0.1% (on ground)
5	Convergence time for relative position knowledge with multiple observers	Convergence observed	2 orbits	1 orbit (in flight)
6	Accuracy of relative position knowledge in presence of maneuvers	Estimate produced	0.1% error relative to range	Estimate produced (in flight)
7	Convergence time for relative position knowledge in presence of maneuvers	Convergence observed	2 orbits	Convergence observed (in flight)
8	Accuracy of absolute position knowledge with multiple observers	Estimate produced	1 km error	850 m error (on ground)
9	Convergence time for absolute position knowledge with multiple observers	Convergence observed	24 hours	Convergence observed (on ground)
10	Time to produce autonomous relative orbit initialization	Initialization produced	2 orbits	5 orbits (in flight)

consists of four identical 6U CubeSats and was launched in July 2023. Prior demonstrations of angles-only navigation in orbit have performed navigation for single observers and single targets only, and have required external orbit information and/or system maneuvers to achieve orbit determination. To overcome these limitations, StarFOX applies the new angles-only Absolute and Relative Trajectory Measurement System (ARTMS). The ARTMS architecture 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 state estimates by seamlessly fusing measurements from multiple observers received over an inter-satellite crosslink within an unscented Kalman filter.

Initial StarFOX experiments have applied ARTMS in flight to successfully perform distributed swarm navigation. Star trackers are used to image other swarm spacecraft and obtain bearing angle measurements, which are optionally shared over the crosslink. Observers are able

to navigate for multiple target spacecraft and display maneuver-free angles-only convergence, achieving relative position estimation errors of 0.5% of target range. Multiple observers are able to cooperate to estimate relative orbits with an improved 0.1% position error relative to range. StarFOX has also provided the first demonstrations of autonomous initialization of relative navigation, in which an unknown target is detected and tracked by an observer and a relative state initialization is computed in flight. Finally, absolute navigation using only inter-satellite bearing angles was investigated, whereupon coarse initial orbit errors were refined from several kilometers to less than 1 km.

StarFOX encountered several challenges in flight, including poorer than expected target visibility and time synchronization issues, both of which degraded navigation performance. It is likely these could have been mitigated pre-launch with improved image simulation capabilities and more extensive hardware-in-the-loop testing. Nevertheless, StarFOX has achieved the majority of its pre-flight goals during its initial experiment period. It therefore presents a significant step towards the autonomy of future distributed missions and more widespread application of vision-based navigation in space.

Acknowledgments

NASA Starling is funded by the Small Spacecraft Technology Program (SSTP) within the NASA Space Technology Mission Directorate. StarFOX development was supported by SSTP cooperative agreement number 80NSSC18M0058 and by the Air Force Office of Scien-

tific Research award number FA9550-21-1-0414.

The authors acknowledge the invaluable contributions of Joshua Sullivan and Adam W. Koenig towards the conception and development of ARTMS and StarFOX. They also thank the NASA Ames Starling team for their tireless support of StarFOX efforts.

References

- [1] Simone D’Amico, Marco Pavone, Shailendhar Saraf, et al. Miniaturized Autonomous Distributed Space System for Future Science and Exploration. In *8th International Workshop on Satellite Constellations and Formation Flying*, Delft, The Netherlands, 2015.
- [2] Byron D Tapley, S Bettadpur, Mo Watkins, and Ch Reigber. The gravity recovery and climate experiment: Mission overview and early results. *Geophysical Research Letters*, 31(9), 2004.
- [3] Gerhard Krieger, Alberto Moreira, Hauke Fiedler, Irena Hajnsek, Marian Werner, Marwan Younis, and Manfred Zink. Tandem-x: A satellite formation for high-resolution sar interferometry. *IEEE Transactions on Geoscience and Remote Sensing*, 45(11):3317–3341, 2007.
- [4] S. D’Amico, J.-S. Ardaens, and R. Larsson. Spaceborne Autonomous Formation-Flying Experiment on the PRISMA Mission. *Journal of Guidance, Control, and Dynamics*, 35(3):834–850, 2012.
- [5] JL Burch, TE Moore, RB Torbert, and BL Giles. Magnetospheric multiscale overview and science objectives. *Space Science Reviews*, 199(1-4):5–21, 2016.
- [6] M. J. Holzinger and M. K. Jah. Challenges and potential in space domain awareness. *Journal of Guidance, Control, and Dynamics*, 41(1):15–18, 2018.
- [7] Daley Arney, Richard Sutherland, John Mulvaney, Devon Steinkoenig, Christopher Stockdale, and Mason Farley. On-orbit Servicing, Assembly, and Manufacturing (OSAM) State of Play. Technical report, NASA, 2021.
- [8] S. Palo, G. Stafford, and A. Hoskins. An Agile Multi-Use Nano Star Camera for Constellation Applications. In *Proceedings of the Small Satellite Conference*, Logan, UT, 2013.
- [9] David C. Woffinden and David K. Geller. Observability Criteria for Angles-Only Navigation. *IEEE Transactions on Aerospace and Electronic Systems*, 45(3):1194–1208, 2009.
- [10] J. Sullivan and S. D’Amico. Adaptive Filtering for Maneuver-Free Angles-Only Navigation in Eccentric Orbits. In *27th AAS/AIAA Space Flight Mechanics Conference*, San Antonio, TX, 2017.
- [11] A. W. Koenig and S. D’Amico. Observability-Aware Numerical Algorithm for Angles-Only Initial Relative Orbit Determination. In *2020 AAS/AIAA Astrodynamics Specialist Conference*, Lake Tahoe, CA, 2020.
- [12] Robert B. Friend. Orbital Express Program Summary and Mission Overview. In Richard T. Howard and Pejmun Motaghedi, editors, *Proceedings of SPIE*, volume 6958, 2008.
- [13] S. D’Amico, J.-S. Ardaens, G. Gaias, H. Benninghoff, B. Schlepp, and J. L. Jørgensen. Non-cooperative Rendezvous Using Angles-Only Optical Navigation: System Design and Flight Results. *Journal of Guidance, Control, and Dynamics*, 36(6):1576–1595, 2013.
- [14] G. Gaias, S. D’Amico, and J.-S. Ardaens. Angles-Only Navigation to a Noncooperative Satellite Using Relative Orbital Elements. *Journal of Guidance, Control, and Dynamics*, 37(2):439–451, 2014.
- [15] Gabriella Gaias and Jean-Sébastien Ardaens. Flight Demonstration of Autonomous Noncooperative Rendezvous in Low Earth Orbit. *Journal of Guidance, Control, and Dynamics*, 41(6):1337–1354, 2017.
- [16] Jean-Sébastien Ardaens and Gabriella Gaias. Angles-Only Relative Orbit Determination in Low Earth Orbit. *Advances in Space Research*, 61(11):2740–2760, 2018.
- [17] H. Sanchez, D. McIntosh, H. Cannon, C. Pires, J. Sullivan, S. D’Amico, and B. O’Connor. Starling-1: Swarm Technology Demonstration. In *32nd Annual Small Satellite Conference*, Logan, UT, 2018.
- [18] Joshua Sullivan, Adam W Koenig, Justin Kruger, and Simone D’Amico. Generalized Angles-Only Navigation Architecture for Autonomous Distributed Space Systems. *Journal of Guidance, Control, and Dynamics*, 44(6):1087–1105, 2021.
- [19] Justin Kruger, Adam W. Koenig, and Simone D’Amico. The Starling Formation-Flying Optical Experiment (StarFOX): System Design, Experiment Design and Pre-Flight Verification. *Journal of Spacecraft and Rockets*, pages 1–23, 2023.

- [20] D. A. Vallado and W. D. McClain. *Fundamentals of Astrodynamics and Applications*. Microcosm Press, Hawthorne, California, 4 edition, 2013.
- [21] S. D’Amico, J.-S. Ardaens, G. Gaias, H. Benninghoff, B. Schlepp, and J. L. Jørgensen. Non-cooperative Rendezvous Using Angles-Only Optical Navigation: System Design and Flight Results. *Journal of Guidance, Control, and Dynamics*, 36(6):1576–1595, 2013.
- [22] S. D’Amico. *Autonomous Formation Flying in Low Earth Orbit*. PhD thesis, Delft University, 2010.
- [23] A. W. Koenig, T. Guffanti, and S. D’Amico. New State Transition Matrices for Relative Motion of Spacecraft Formations in Perturbed Orbits. *Journal of Guidance, Control, and Dynamics*, 40(7):1749–1768, 2017.
- [24] Tommaso Guffanti and Simone D’Amico. Linear Models for Spacecraft Relative Motion Perturbed by Solar Radiation Pressure. *Journal of Guidance, Control, and Dynamics*, 42(9):1962–1981, 2019.
- [25] J. Sullivan, T. A. Lovell, and S. D’Amico. Angles-Only Navigation for Autonomous On-Orbit Space Situational Awareness Applications. In *AAS/AIAA Astrodynamics Specialist Conference*, Snowbird, UT, 2018.
- [26] Justin Kruger, Katherine Wallace, Adam W. Koenig, and Simone D’Amico. Autonomous Angles-Only Navigation for Spacecraft Swarms around Planetary Bodies. In *2021 IEEE Aerospace Conference*, Big Sky, MT, 2021.
- [27] K. Iiyama, J. Kruger, and S. D’Amico. Autonomous Distributed Angles-Only Navigation and Timekeeping in Lunar Orbit. In *ION International Technical Meeting*, Long Beach, CA, 2022.
- [28] Justin Kruger and Simone D’Amico. Observability Analysis and Optimization for Angles-Only Navigation of Distributed Space Systems. In *11th International Workshop on Satellite Constellations and Formation Flying*, Milano, Italy, 2022.
- [29] K. T. Alfriend, editor. *Spacecraft Formation Flying: Dynamics, Control, and Navigation*. Butterworth-Heinemann/Elsevier, 2010.
- [30] Oliver Montenbruck and Eberhard Gill. *Satellite Orbits: Models, Methods and Applications*. Springer, 2012.
- [31] Joshua Sullivan, A. Koenig, and Simone D’Amico. Improved Maneuver-Free Approach to Angles-Only Navigation for Space Rendezvous. In *26th AAS/AIAA Space Flight Mechanics Meeting*, Napa, CA, 2016.
- [32] J. Sullivan and S. D’Amico. Nonlinear Kalman Filtering for Improved Angles-Only Navigation Using Relative Orbital Elements. *Journal of Guidance, Control, and Dynamics*, 40(9):2183–2200, 2017.
- [33] Justin Kruger and Simone D’Amico. Autonomous angles-only multitarget tracking for spacecraft swarms. *Acta Astronautica*, 189:514–529, 2021.
- [34] Ba-ngu Vo, Mahendra Mallick, Yaakov Barshalom, et al. Multitarget Tracking. In J. G. Webster, editor, *Wiley Encyclopedia of Electrical and Electronics Engineering*. Springer International Publishing, 2015.
- [35] Xiaowei Wan, Gangyi Wang, Xinguo Wei, Jian Li, and Guangjun Zhang. Star Centroiding Based on Fast Gaussian Fitting for Star Sensors. *Sensors*, 18(9), 2018.
- [36] D. Mortari, M. A. Samaan, C. Bruccoleri, and J. L. Junkins. The Pyramid Star Identification Technique. *Navigation*, 51(3):171–183, 2004.
- [37] J. R. Wertz. *Spacecraft Attitude Determination and Control*. Springer Science & Business Media, 2012.
- [38] S. Blackman. Multiple hypothesis tracking for multiple target tracking. *IEEE Aerospace and Electronic Systems Magazine*, 19(1):5–18, 2004.
- [39] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*, page 226–231, 1996.
- [40] Nathan Stacey and Simone D’Amico. Autonomous swarming for simultaneous navigation and asteroid characterization. In *AAS/AIAA Astrodynamics Specialist Conference*. Snowbird, UT, 2018.
- [41] Yunpeng Hu, Inna Sharf, and Lei Chen. Distributed orbit determination and observability analysis for satellite constellations with angles-only measurements. *Automatica*, 129:109626, 2021.
- [42] C. Beierle and S. D’Amico. Variable magnification optical stimulator for training and validation of spaceborne vision-based navigation. *Journal of Spacecraft and Rockets*, 56(4):1060–1072, 2019.