Towards Autonomous Space Situational Awareness at the Edge: Flight Results from NASA Starling

Justin Kruger, Alex Mescher, Aviad Golan, Sumant Sharma, and Simone D'Amico
EraDrive, Inc.

ABSTRACT

The large number of satellites carrying onboard star trackers encourages their usage as opportunistic data collection platforms for space situational awareness (SSA), with the potential to collect more timely, more persistent, and more accurate data than ground-based assets. It is especially valuable if SSA data is processed onboard to improve autonomy and safety in real time. EraDrive proposes a novel architecture in which distributed networks of space-based observers enable autonomous space traffic coordination (STC) at both the global and individual satellite level. The architecture consists of online 'Era-Node' optical sensing and processing modules, and an offline 'Era-Cloud' procedure which maintains a global catalog of Resident Space Objects (RSO). The online process detects RSOs in images acquired by the Era-Node camera and performs identification, tracking, and orbit refinement. Refined state estimates are used to enable autonomous collision avoidance for the observer. Data is shared between modules over a crosslink to enable distributed SSA and STC. The Era-Cloud process fuses space-based data from proliferated Era-Node modules with external ground-based data sources to produce an authoritative RSO catalog. Era-Cloud products are uplinked to Era-Node modules and provided to operators to reduce conjunction assessment uncertainties.

Feasibility and performance of the architecture are assessed through a combination of flight experiments for individual Era-Nodes, aboard the NASA Starling mission, and numerical simulations of their proliferation. In flight, multiple Era-Node modules performed autonomous SSA, matching measurements of 3000 unique RSOs to catalog identities with mean residuals of less than 100 m. An onboard attitude planner was able to choose RSOs of interest for observation and maintain custody by centering targets in the field of view. Measurements were subsequently used to refine stale catalog orbit information on board, reducing position errors by 50% compared to two-line elements. Measurements were also shared between multiple Era-Nodes to enable initialization of new RSO orbit estimates. Era-Cloud performance is characterized through simulations of the complete system, calibrated using Starling flight data. Proliferated Era-Node modules achieved 95% coverage of the LEO RSO population in two hours, with median revisit times of <5 minutes and orbit errors of tens of meters, superior to typical SSA products. The proposed architecture therefore provides multilayered benefits to the safety and autonomy of individual satellites as well as space sustainability at scale.

1. INTRODUCTION

The rise of dynamic space operations at scale has made timely, persistent, and accurate space situational awareness (SSA) an increasingly critical capability. Plans for upcoming megaconstellations propose millions of satellites to be launched over the next several decades [1], and there is growing demand for more agile satellite behavior to fulfill objectives related to on-orbit logistics and co-orbital warfare [2]. Robust and accessible SSA products are needed to ensure safe interactions between large numbers of potentially maneuvering, autonomous, or non-cooperative assets.

SSA has typically been performed using ground-based resources for both sensing and compute, in which measurements of resident space objects (RSOs) are processed to produce regularly-updated RSO catalogs. However, this process is characterized by inherent limitations. Ground-based sensing places severe constraints on the sizes of objects that can be reliably observed, and limits the potential quality and coverage of measurements [3]. Furthermore, the restriction to centralized ground-based processing guarantees significant delays between the taking of a measurement, the updating of the catalog, and the ability of a satellite to use this information for collision avoidance.

Dedicated space-based SSA initiatives such as the NorthStar constellation [4] and the United States' Geosynchronous SSA Program partially address the first limitation, but their proliferation and broader coverage remain limited. Given that the majority of modern satellites carry star trackers or other cameras, an alternative paradigm is to leverage existing satellites as opportunistic platforms to passively image RSOs and collect SSA data. SpaceX's Stargaze program states that data from 30,000 star trackers will be used to detect 30 million RSO transits daily across its Starlink fleet [5]. An example star tracker image, obtained by EraDrive on the NASA Starling mission, is shown in Figure 1.

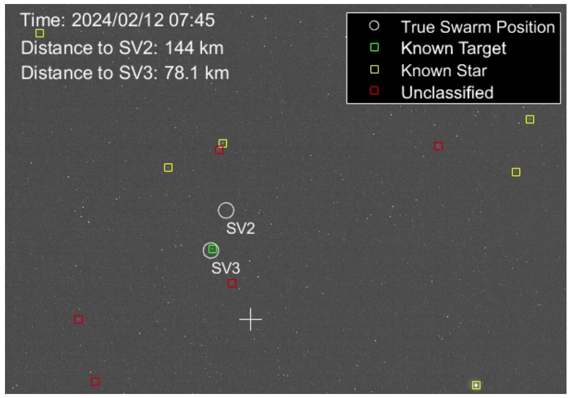

Fig. 1: Star tracker imagery from NASA's Starling swarm. SV3 is an RSO CubeSat target.

Data-oriented delays remain present because Stargaze (and similar space-based SSA platforms) prioritize downlinking sensor data to be processed on the ground, rather than performing SSA online. For dynamic space traffic coordination (STC), it is instead valuable for satellites to process and act upon SSA data in flight, such that they are aware of their environments and can conduct safe decision-making and maneuver planning in real time. However, a single satellite will naturally possess less measurement information than a centralized on-ground database when generating SSA products, and on-orbit computational resources are constrained by comparison.

In response, EraDrive has proposed a new paradigm for SSA and STC which fuses the benefits of on-ground and on-orbit autonomy. On-orbit sensors are used to collect measurements of RSOs, which are processed online by the host satellite in order to provide real-time SSA knowledge. On-orbit crosslinks enable measurement sharing, handshaking, and cooperative maneuver planning for collision avoidance between satellites. Measurements from all on-orbit sensors are also downlinked and combined to generate a global RSO catalog (in cooperation with other data sources). Instead of requiring a dedicated constellation, on-orbit measurements are collected in an opportunistic manner by existing cameras present on spacecraft (e.g. star trackers), thus turning every satellite into an SSA platform without being locked to a single operator. Furthermore, crosslinks enable responsive STC at the edge through a virtual constellation of interconnected modules. The resulting distributed SSA/STC network requires minimal additional resources while simultaneously improving coverage, accuracy, and operational autonomy via space-based sensing.

Feasibility of the EraDrive concept is verified through two avenues: flight experiments performing real-time SSA at the edge, and simulation experiments characterizing coverage and accuracy of the distributed SSA network. Experiments in 2025 and 2026 aboard the NASA Starling swarm have successfully demonstrated key elements of the proposed architecture, including RSO identification, orbit refinement, orbit initialization, and persistent custody. Scaling of these outcomes to hundreds of distributed observer satellites and tens of thousands of RSOs is explored through high-fidelity simulations and analysis of SSA performance metrics. Together, these provide the first demonstrations of fully autonomous SSA on orbit and the first in-depth studies of proliferated, opportunistic space-based SSA in literature.

After this introduction, relevant mathematical background is presented. Next, components of the proposed SSA/STC architecture are outlined in detail with reference to the current state of the art. The NASA Starling mission is then described, followed by flight results for novel autonomous SSA capabilities at the edge. Finally, the simulation environment and overall SSA network performance are characterized, followed by conclusions.

2. PRELIMINARIES

2.1 Measurement Models

It is assumed target RSOs are unresolved point sources in camera images and can be reduced to time-tagged bearing angle pairs. Two rotating coordinate frames are defined. First, consider the Radial/Tangential/Normal (RTN) frame of a space-based observer, denoted $\mathscr{R}$. It is centered on and rotates with the observer and consists of orthogonal basis vectors $\hat{\boldsymbol{x}}^{\mathscr{R}}$ (directed along the observer's absolute position vector); $\hat{\boldsymbol{z}}^{\mathscr{R}}$ (directed along the observer's orbital angular momentum vector); and $\hat{\boldsymbol{y}}^{\mathscr{R}} = \hat{\boldsymbol{z}}^{\mathscr{R}} \times \hat{\boldsymbol{x}}^{\mathscr{R}}$ [6]. Similarly, define a frame $\mathscr{W}$ using $\hat{\boldsymbol{y}}^{\mathscr{W}}$ (directed along the observer's velocity vector); $\hat{\boldsymbol{z}}^{\mathscr{W}} = \hat{\boldsymbol{z}}^{\mathscr{R}}$; and $\hat{\boldsymbol{x}}^{\mathscr{W}} = \hat{\boldsymbol{y}}^{\mathscr{W}} \times \hat{\boldsymbol{z}}^{\mathscr{W}}$. $\mathscr{W}$ only differs from $\mathscr{R}$ by a rotation of the observer flight path angle ϕ_f about $\hat{\boldsymbol{z}}^{\mathscr{R}}$ with $\phi_f \approx 0$ in near-circular orbits [6]. Finally, define the observer vision-based sensor (VBS) coordinate frame $\mathscr{V}$ consisting of orthogonal basis vectors $\hat{\boldsymbol{x}}^{\mathscr{V}}, \hat{\boldsymbol{y}}^{\mathscr{V}}, \hat{\boldsymbol{z}}^{\mathscr{V}}$ where $\hat{\boldsymbol{z}}^{\mathscr{V}} = \hat{\boldsymbol{x}}^{\mathscr{V}} \times \hat{\boldsymbol{y}}^{\mathscr{V}}$ is aligned with the camera boresight. The VBS may be pointed as necessary to keep targets in the field of view (FOV). Figure 2 illustrates this scenario, where the camera boresight is pointed in the anti-velocity direction for simplicity.

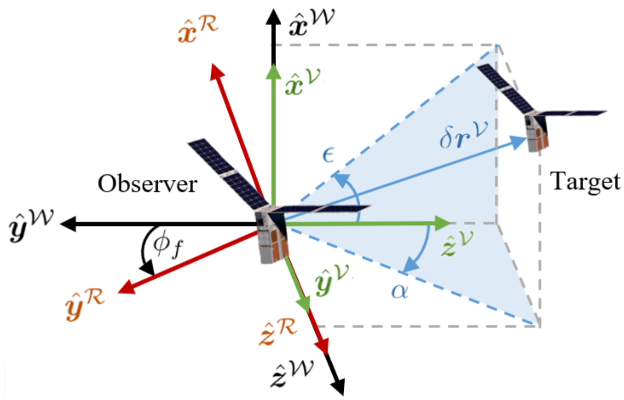

Fig. 2: Definition of coordinate frames and bearing angles with VBS pointing in the $-\hat{\boldsymbol{y}}^{\mathscr{W}}$ direction.

Bearing angles consist of azimuth and elevation $[\alpha, \varepsilon]^\top$. Geometrically, they subtend the line-of-sight vector $\delta\boldsymbol{r}^{\mathscr{V}}$ from observer to target, and are an analog for the image frame coordinates of the target object pixel cluster. In this work, the measurement model $\boldsymbol{y}$ for bearing angles from observer o to target t is [7]

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

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

Measurements and states are also referenced with respect to an inertial reference frame centered on an arbitrary central body, denoted $\mathscr{I}$. This work primarily applies the Earth-centered inertial (ECI) J2000 frame. Camera measurements are related to $\mathscr{I}$ as per $\delta\boldsymbol{r}^{\mathscr{I}} = {}^{\mathscr{V}}\overrightarrow{\boldsymbol{R}}^{\mathscr{I}} \delta\boldsymbol{r}^{\mathscr{V}}$, where ${}^{\mathscr{V}}\overrightarrow{\boldsymbol{R}}^{\mathscr{I}}$ denotes a rotation from frame $\mathscr{V}$ into frame $\mathscr{I}$. This rotation is generally computed by performing attitude determination using stars identified by the VBS [8]. Other rotations ${}^{\mathscr{R}}\overrightarrow{\boldsymbol{R}}^{\mathscr{I}}$ and ${}^{\mathscr{W}}\overrightarrow{\boldsymbol{R}}^{\mathscr{I}}$ can be computed using the observer's absolute orbit estimate, as derived from GNSS, or through alternative means such as optical measurements of RSOs from the onboard camera [9]. The true equator mean equinox (TEME) frame $\mathscr{T}$ is employed when working with two-line elements and is related to $\mathscr{I}$ through nutation and precession terms which vary with the current epoch [10].

2.2 Dynamics Models

The use of space-based cameras to collect passive optical measurements is a form of angles-only navigation, which is known to be characterized by weak observability [11, 12]. Furthermore, tracking (or simulating) an arbitrarily large RSO population using limited hardware implies computational efficiency is needed. Proper selection of a state parameterization and dynamics model is therefore crucial to maximize SSA accuracy and robustness.

To represent the states of observer and target objects, quasi-nonsingular absolute Orbit Elements (OE) and Relative Orbit Elements (ROE) are leveraged. The absolute orbit $\boldsymbol{\alpha}$ is defined as

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

where $a, e, i, \Omega, \omega, M$ are the Keplerian orbit elements. The relative orbit of each target detected by the onboard sensor, denoted $\delta\boldsymbol{\alpha}$, is described by the quasi-nonsingular ROE adopted by D'Amico [13]. Each of these ROE is function of the orbit elements of t and o, given by

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

where $[\delta e_x, \delta e_y]$ are components of the relative eccentricity vector with phase ϕ and $[\delta i_x, \delta i_y]$ are components of the relative inclination vector with phase θ. Fully nonsingular OE and ROE have also been developed for equatorial orbits.

When high fidelity is required, the absolute OE are propagated using fourth-order Runge-Kutta integration of Gauss' Variational Equations (GVE). For state $\boldsymbol{\alpha}$, the osculating OE evolve according to

$$\dot{\boldsymbol{\alpha}} = G(\boldsymbol{\alpha})\boldsymbol{d}^{\mathscr{R}} \tag{5}$$

where $G \in \mathbb{R}^{6\times3}$ is the GVE matrix [14] and $\boldsymbol{d}^{\mathscr{R}} \in \mathbb{R}^3$ is the perturbing acceleration expressed in $\mathscr{R}$. Depending on the orbit regime, common perturbations include spherical harmonic gravity terms, atmospheric drag, third-body gravity and solar radiation pressure [6]. Importantly, the OE and ROE vary slowly with time, which enables accurate, efficient numerical integration of the GVE with large time steps [15]. Accurate analytical dynamics models have also been developed for when computational efficiency is paramount [16, 17].

Two-line elements (TLEs) are also commonly used to parametrize object states for SSA [10]. This work leverages TLEs as a compact state representation which contains additional information (e.g. an object identification number, a drag term, and rates of change of mean motion) compared to standard six-dimensional OE vectors. The analytic SGP4 and SDP4 models are used to propagate TLE states en masse, which incorporate coarse time-averaged perturbation effects without excessive computational cost. However, TLEs do not include uncertainties by default.

3. ARCHITECTURE

3.1 State of the Art

The majority of current SSA architectures integrate ground-based sensors, which gather RSO measurements, with ground-based software, to produce catalogs of RSO orbit information. Sensors include radar stations and optical telescopes with varying fields of view (FOV), detection ranges, and availability [18]. RSO catalogs fusing these measurements are produced by various government and commercial entities such as the United States Space Force (USSF), the European Union, LeoLabs, COMSPOC, ExoAnalytic, Slingshot Aerospace, and Kayhan Space, with differing levels of detail, timeliness, coverage, and accuracy [19, 20, 21, 22]. Initiatives such as the Traffic Coordination System for Space (TraCSS) are also attempting to fuse data across these sources to improve catalog quality [23]. Catalog information is employed downstream to facilitate target tracking, conjunction assessment and collision avoidance.

As of July 2026, the freely available, authoritative Space-Track catalog contains 51,200 objects, most of which are $\geq$10 cm in diameter [24]. The approximate level of position error in Space-Track TLEs is 5-100 m in the radial and cross-track directions and 1-3 km per day in the along-track direction [25, 26]. Mean revisit rates are on the order of 10 hours with an update latency of 8 hours. Significant attention has been paid to improving upon Space-Track by integrating additional sensors or generating additional products such as TLE uncertainties [27]. However, the number of available sensors and their quality of data are physically limited, meaning significant improvements to revisit rates and orbit errors can only be achieved for subsets of tracked objects. In general, current SSA architectures are considered outdated with large state uncertainties and poor latency, which reduces the efficiency and safety of STC [3, 28].

A newer paradigm for SSA is the integration of space-based sensing. In comparison to ground-based sensing, space-based sensors have the potential to provide more consistent measurement availability, more flexible pointing, and physically smaller inter-object distances, enabling tracking of small debris objects [29]. Examples include the USSF Space Surveillance Network, which incorporates multiple space-based assets [30] and NorthStar, which has launched multiple satellites as part of a dedicated SSA constellation [4]. However, these systems focus on narrow-FOV, high-sensitivity optical payloads which restricts overall proliferation and coverage. To overcome this, there is growing interest in leveraging star tracker data from commercial satellites to gather RSO observations at scale. Although star trackers are less sensitive than a dedicated telescope, they generally feature a wider FOV and are present on almost every modern satellite. SpaceX's Stargaze program intends to aggregate star tracker data from tens of thousands of Starlink satellites [5] and recent work by the authors has successfully applied star trackers for autonomous optical navigation of distributed space systems (DSS) [31]. However, the achievable performance of this concept when proliferated has not yet been verified in literature and most tests have been restricted to ground-based trials [32, 33].

Finally, recall that an underlying motivation for robust SSA is safe, efficient STC. Current space-based SSA systems rely on downlinking measurements and do not use or update RSO catalogs on board. The capacity for involved satellites to use this SSA information for reactive, autonomous coordination in the presence of a growing space population is therefore limited. Recent experiments aboard the NASA Starling mission have explored the potential for autonomous SSA on orbit [34] but no satellite has yet implemented a complete SSA/STC pipeline in flight.

3.2 New Architecture

The EraDrive SSA architecture proposes a novel capability in which distributed networks of space-based observers enable autonomous STC, at both the global and individual satellite level. Conceptually, the architecture consists of two layers: online 'Era-Node' optical sensing, processing, and communication modules, and an offline 'Era-Cloud' data aggregation procedure which maintains a global catalog of RSOs. Each Era-Node possesses an onboard camera with which to collect optical angles-only measurements of target objects, and an inter-satellite link (ISL) for data sharing and traffic coordination with other observers. Observers do not necessarily share operators or mission objectives but instead conduct SSA/STC as background processes as members of a virtual constellation. RSOs which are not carrying Era-Node modules are considered non-cooperative. Observer SSA data is downlinked regularly and provided to operators and to a centralized database to enable the Era-Cloud layer on the ground. A notional scenario is illustrated in Figure 3.

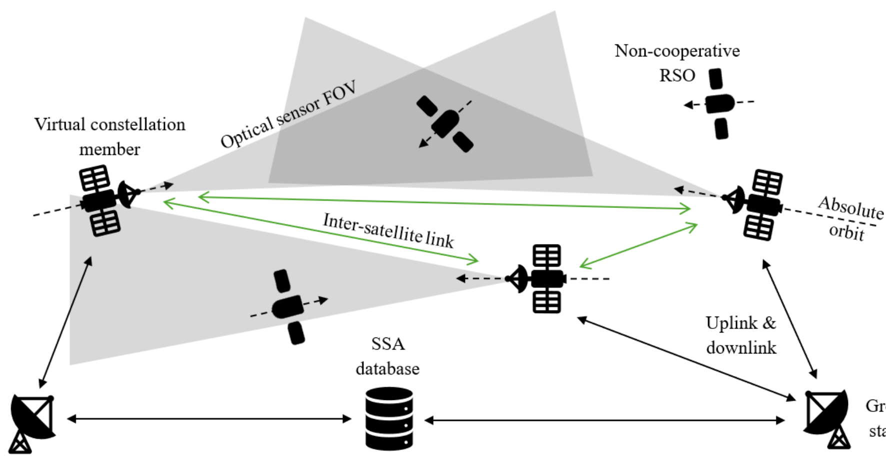

Fig. 3: Notional illustration of an SSA/STC architecture scenario.

The online process runs in a decentralized, distributed fashion aboard the N spacecraft in the virtual constellation. Each Era-Node detects RSOs in images acquired by its camera and compares bearing angles with predicted locations from the RSO catalog on board. Upon successful identity match, local bearing angles are fused with other available sources (e.g. onboard GNSS, or bearing angles broadcast over the ISL by other Era-Node modules) to produce refined state estimates for the observer and measured RSOs concurrently. If the RSO is unmatched, a batch of measurements is collected over time to compute an on-orbit initialization. Refined state estimates are used to perform autonomous conjunction assessment and collision avoidance maneuver planning for the observer, if operationally allowed. Intermediate SSA/STC products (e.g. bearing angles, catalog orbit updates, and planned maneuvers) are shared with the ground and with other Era-Nodes over a radio frequency (RF) network. Nominal onboard update rates range from 0.1-1 Hz depending on available compute. Downlink nominally occurs several times per day, noting that Era-Node data products are much smaller than raw images.

The Era-Cloud process runs in a centralized fashion on the ground. Downlinked Era-Node data products from the virtual constellation are fused with available measurements and RSO catalog updates from external third-party providers.

Subsequently, the Era-Cloud procedure outputs an updated RSO catalog considering all sources of measurement information. A nominal catalog update rate is 1 hour. The resulting products are uplinked to Era-Nodes at a regular cadence, and may also be accessed by general operators to reduce conjunction assessment uncertainties. Uplink may be daily for some active satellites or at a reduced cadence (e.g. weekly) depending on the available bandwidth. A block diagram of the proposed SSA/STC architecture is presented in Figure 4.

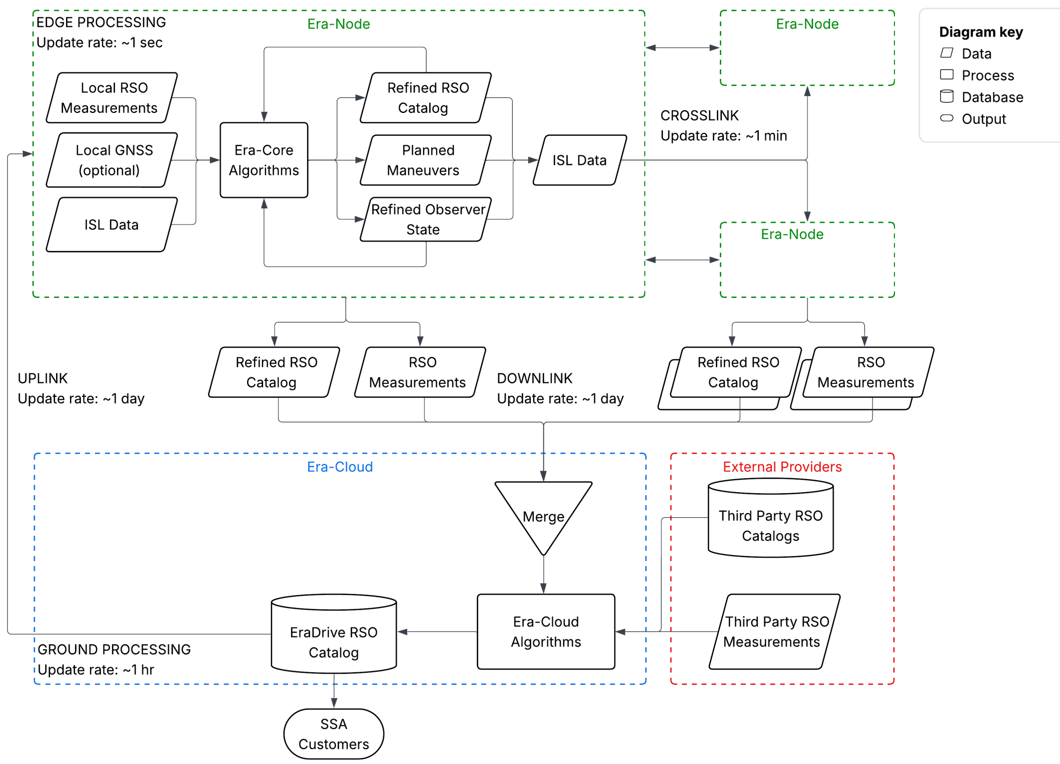

Fig. 4: High-level block diagram of the proposed SSA architecture.

When considered in the context of current space-based or ground-based SSA architectures, the EraDrive concept proposes several key enhancements. First, the use of proliferated star tracker data, agnostic to any specific operator, provides the highest possible levels of measurement availability, coverage, and revisit rate. When aggregated on the ground and fused with other sources, the resulting multi-observer data is expected to enable higher accuracies and lower uncertainties for SSA products. Second, the inclusion of edge processing enables onboard autonomy and reduces latency significantly. Individual satellites are given visual awareness of their surroundings and access to more accurate SSA knowledge, meaning that maneuver planning and STC can be performed online. The ISL also enables creation of a virtual constellation which is able to share data and coordinate maneuvers. The following sections confirm the viability of the architecture and these benefits through complementary flight experiments and simulation studies.

4. STARLING MISSION

Feasibility of the online component is verified through flight experiments aboard the NASA Starling swarm. Starling was proposed as a testbed for autonomous swarming technologies and consists of four propulsive 6U CubeSats, launched into approximately sun-synchronous low Earth orbit (LEO) in July 2023. Its primary mission successfully concluded in May 2024 [35], which has been followed by multiple mission extensions implementing novel autonomy capabilities. EraDrive's flight software and flight experiment contributions have encompassed angles-only absolute and relative orbit determination for the swarm [31], autonomous space-based SSA [34], autonomous optical positioning, navigation and timing (PNT) using RSO measurements [9], autonomous attitude planning, and autonomous maneuver planning for swarm maintenance and reconfiguration. The swarm remains in operation as of July 2026.

4.1 Flight Hardware

For optical navigation and SSA, each CubeSat carries two Blue Canyon Technologies Nano Star Trackers (NST) aligned in antiparallel directions, either of which may collect imagery [36]. NST specifications 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 and passing RSOs. Images are provided to the Era-Node flight software with a nominal cadence of 60 seconds. The other NST provides a backup attitude solution using visible stars.

Table 1: Imaging parameters of the NST.

Parameter	Value
Image Size (pixels)	1280×1024
FOV (°)	12×10
Pixel Depth	10-bit (grayscale)
Max. Detectable Magnitude	7.5

The software also receives inputs from the onboard GPS receiver, which optionally provides position/velocity/time (PVT) solutions, and from a swarm ISL, via 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 is a Xiphos Q7S running at approximately 700 MHz. Each satellite also possesses a cold gas propulsion system with four separate nozzles, used to conduct swarm phasing and reconfiguration maneuvers.

Figure 5 shows the four Starling spacecraft during integration at the NASA Ames Research Center, together with an artist's visualization of the swarm in flight. The spacecraft are referred to as SV1, SV2, SV3 and SV4 or more affectionately as Blinky, Pinky, Inky, and Clyde.

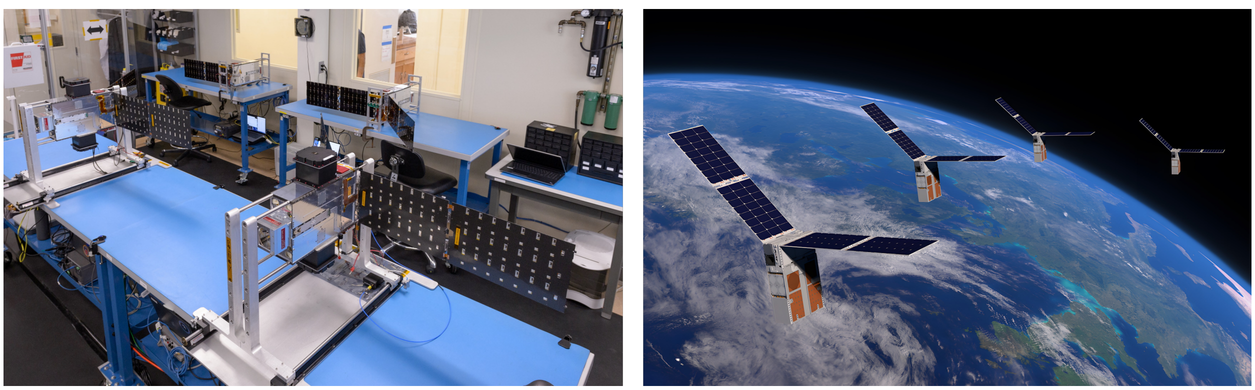

Fig. 5: The four Starling CubeSats during integration (left) (credit: NASA / Dominic Hart) and an artist's impression of the Starling swarm (right) (credit: NASA / Blue Canyon Technologies).

4.2 Flight Software

Starling applies EraDrive's 'Era-Core' flight software to perform online SSA. Era-Core is a modular guidance, navigation and control (GNC) architecture that provides autonomous angles-only navigation for DSS orbiting an arbitrary central body, without reliance on maneuvers or external measurement sources (among other GNC capabilities). Many of its constituent algorithms were originally developed as part of the Starling Formation-Flying Optical Experiment (StarFOX), which was the first on-orbit demonstration of autonomous vision-based navigation for a satellite swarm [31, 37]. The StarFOX algorithms have been augmented with new Era-Core systems which enable tracking of RSOs outside the Starling swarm, including their identification, custody, state initialization, and state refinement. Together, Era-Core, the NST and the ISL make up the Era-Node module on each Starling satellite.

A high-level overview of Era-Core, as implemented aboard Starling, is shown in Figure 6. The flight software consists of four core modules, in green: Image Processing (IMP), Batch Orbit Determination (BOD), Sequential Orbit Determination (SOD), and Maneuver Planning (MAP). Each active observer in the DSS runs an identical Era-Core instance to navigate for itself, other observers, and any detected RSOs. A typical sequence of execution is as follows.

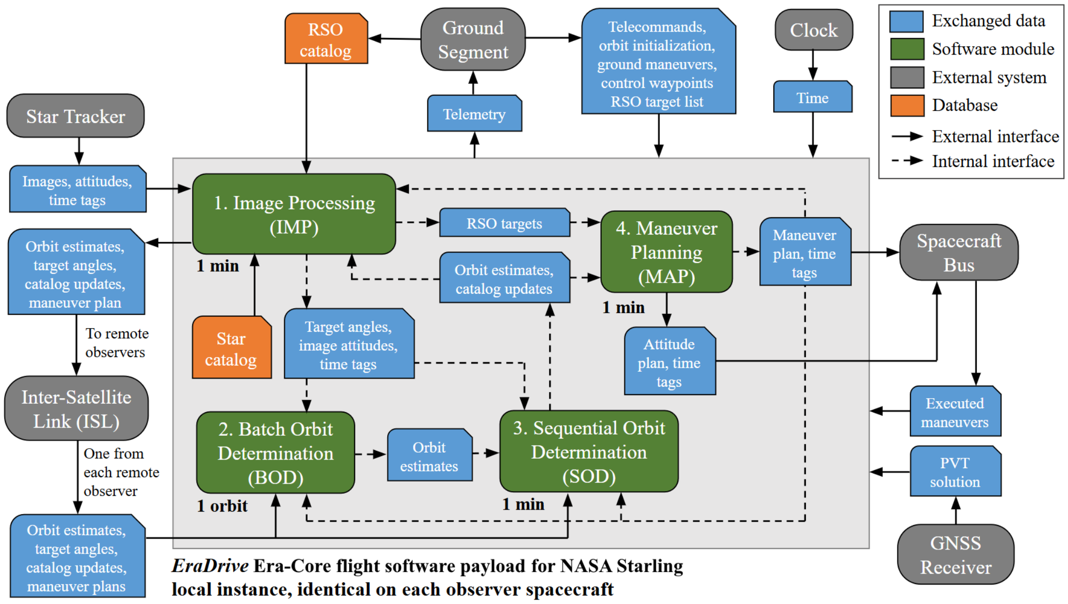

Fig. 6: High-level block diagram of the Era-Core autonomous SSA flight software for Starling.

First, IMP processes the new image from the onboard star tracker to produce target measurements. Bearing angles to visible DSS targets and RSO targets are collected. A star catalog is used to identify stars for attitude determination. An uplinked RSO catalog (consisting of TLEs and orbit uncertainties, if available) is used to identify RSOs. If orbit estimates from SOD are available, they are leveraged to identify DSS targets. IMP sends computed bearing angles and attitudes to BOD and SOD to enable orbit estimation. It also prepares the outbound ISL message by collating onboard bearing angle measurements, orbit estimates, catalog updates, and maneuver plans, to be broadcast to other observers for distributed coordination. IMP executes every 60 seconds.

BOD uses the batch of local bearing angles from IMP, as well as batches of remote bearing angles from the ISL, to compute target orbit initializations. The BOD output is sent to SOD to either initialize navigation or as a consistency check if SOD is already initialized. BOD executes once per orbit.

SOD seamlessly fuses new bearing angles from IMP and remote bearing angles from the ISL to refine target orbit estimates. RSO measurements may be used to refine the observer's absolute orbit for resilient positioning, or to refine RSO orbit information in the onboard catalog. GNSS is optionally used to make the observer orbit estimate more robust. Orbit estimates and catalog updates are provided to IMP (to aid target tracking) and MAP (to enable maneuver planning). SOD executes every 60 seconds.

Finally, MAP leverages a suite of control solvers to produce maneuver plans for STC. Updated DSS orbit estimates from SOD are compared to control waypoints from the ground, which define desired DSS geometry. If the current control error exceeds a user-specified tolerance, a maneuver plan is computed. The minimum radial-normal separation between DSS members is also predicted over a time horizon. If this separation violates a user-specified bound, an escape maneuver is computed to introduce passive safety. Similarly, the probability of collision between the observer and each cataloged RSO is computed over a time horizon. If the RSO miss distance violates a user-specified bound, a collision avoidance plan is computed. Maneuver plans are provided to the spacecraft bus (for execution), to BOD and SOD (for orbit estimation), and to IMP (for ISL broadcast). MAP executes every 60 seconds to check safety conditions, although the cadence at which new maneuvers are computed and executed may vary.

If the ground has provided a list of desired RSO targets to observe, IMP uses the onboard catalog to determine which of these RSOs (if any) can be observed in the next measurement epoch. The best observation candidate is chosen based on user-specified criteria. MAP then uses orbit estimates from SOD, as well as the catalog information for the desired target, to compute attitude commands that center the target in the field of view. These commands are provided to the spacecraft bus. All core software modules receive telecommands, orbit information, maneuver plans, and control waypoints from the ground, and produce telemetry which is downlinked to the ground.

4.3 Flight Experiments

Multiple experiment objectives were defined for the Starling mission follow-on phase, focused on demonstrating key capabilities for autonomous SSA. Table 2 presents experiment requirements. Each experiment is discussed in a dedicated algorithm and results subsection in Sections 5 and 6, respectively.

Table 2: SSA experiment objectives and requirements.

Experiment	Requirement	Success Criteria
RSO Identification	Era-Core shall assign identities to measurements of RSOs outside the swarm that appear in star tracker images.	Match an unidentified visible object to an ID number in the onboard RSO catalog using Era-Core algorithms in flight.
RSO Custody	Era-Core shall maintain visual custody of specified RSOs outside the swarm using the onboard star trackers.	Image a target RSO using an attitude profile computed by Era-Core algorithms in flight, outside the default FOV and attitude.
RSO Refinement	Era-Core shall refine the orbit states of RSOs in the onboard catalog using new measurements from star tracker images.	Demonstrate an improvement in orbit state knowledge for an RSO refined on board, in comparison to existing unrefined TLEs.
RSO Initialization	Era-Core shall generate initial relative orbit estimates for target RSOs using short arcs of bearing angles.	Generate an initial orbit estimate in flight using measurements spanning less than one orbit period.

Approximately 48-72 hours before each flight experiment, a set of uplink products was prepared by EraDrive, encompassing a telecommand table, a-priori maneuver plans, control waypoints, and an RSO catalog. The telecommand table included algorithm parameters, mode switches, and orbit initializations for each swarm spacecraft. Control waypoints specified a desired relative orbit configuration for the swarm. The RSO catalog featured a subset of TLEs from the Space-Track database at the preparation epoch, filtered by user-specified criteria (e.g. restricted to specific altitudes, epochs, or constellations). The lag time between experiment preparation and execution meant catalog information was 2-4 days old before being used on board. The Era-Core software then operated during discrete experiment periods lasting 24-72 hours each. The default attitude aligned the chosen star tracker in the velocity or anti-velocity direction, depending on the experiment. After completion, Era-Core telemetry was downlinked for analysis. Space-Track data and operator ephemerides (e.g. via the SpaceX Starlink API) were used to assess the accuracy of SSA products. Table 3 provides an example of Starling swarm orbits, noting that the swarm's absolute and relative orbits evolve over time.

Table 3: Mean quasi-nonsingular OE and ROE of the Starling swarm.

Satellite	Formation as of 12/06/25 00:00:00 UTC
SV4	$\boldsymbol{\alpha}$ = [6925 km, 0.0007, 0.0014, 99.38°, -52.40°, 43.11°]
SV3	Not included in experiments
SV2	$\delta\boldsymbol{\alpha}$ = [-4, -135890, -1090, 199, 0, 2853] m
SV1	$\delta\boldsymbol{\alpha}$ = [-33, -55624, -8323, -1594, 604, 63306] m

5. FLIGHT ALGORITHMS

5.1 RSO Identification

RSO identification is performed using the onboard RSO catalog within IMP. When a new image arrives at epoch t_{im}, all catalog orbits are propagated from TLE epoch t_{cat} to t_{im} using SGP4. During experiments, the onboard catalog contained up to 25,000 objects, all of which were efficiently propagated for each new image (i.e. every 60 seconds).

IMP then uses the propagated orbits to compute modeled measurement $\boldsymbol{y}_i^{\mathscr{V}}$ for the i'th RSO. If $\boldsymbol{y}_i^{\mathscr{V}}$ lies within the FOV, and the modeled inter-satellite distance (ISD) between the observer and target $||\delta\boldsymbol{r}^{\mathscr{V}}||_2$ is less than a user-specified maximum, its corresponding bearing angle covariance $\boldsymbol{P}_i^{\mathscr{V}}$ is computed. This covariance describes the measurement uncertainty due to TLE orbit uncertainty. Various methods have been proposed for computing this uncertainty (which is typically not provided by Space-Track), such as generation of error samples from the time evolution of TLE data [38]. For simplicity, this work applies an empirical model derived from Starling flight data [39]. An ROE uncertainty vector is defined for the RSO orbit with respect to the observer's orbit, which is converted to $\boldsymbol{P}_i^{\mathscr{V}}$ in bearing angle space via an unscented transform. An example uncertainty vector is

$$\boldsymbol{\sigma}_{\delta\boldsymbol{\alpha}} = [100, 500, 100, 100, 100, 100]^\top + t_{\text{prop}}[50, 2000, 50, 50, 50, 50]^\top \text{m} \tag{6}$$

for t_{prop} expressed in days. This model aims to capture the time dependence and shape dependence of TLE uncertainties (which are typically largest in the velocity direction and grow with time) in relation to the image plane.

Let σ_{ij} denote the Mahalanobis distance between the predicted bearing angles $\boldsymbol{y}_i^{\mathscr{V}}$ of the i'th cataloged RSO and the measured bearing angles $\boldsymbol{y}_j^{\mathscr{V}}$ of the j'th non-stellar object in the image, with

$$\sigma_{ij} = \sqrt{(\boldsymbol{y}_i^{\mathscr{V}} - \boldsymbol{y}_j^{\mathscr{V}})^\top (\boldsymbol{P}_i^{\mathscr{V}})^{-1} (\boldsymbol{y}_i^{\mathscr{V}} - \boldsymbol{y}_j^{\mathscr{V}})} \tag{7}$$

To minimize erroneous assignments, the i'th RSO is only matched to the j'th detected object if

$$\sigma_{ij} \leq \varepsilon_{\text{assign}} \quad ; \quad \sigma_{kj} \geq \varepsilon_{\text{ambig}} \; \forall \, k \neq i \quad ; \quad \sigma_{il} \geq \varepsilon_{\text{ambig}} \; \forall \, l \neq j \tag{8}$$

where $\varepsilon_{\text{assign}}$ and $\varepsilon_{\text{ambig}}$ are user-specified parameters satisfying $\varepsilon_{\text{ambig}} > \varepsilon_{\text{assign}} > 0$. The angular distance $||\boldsymbol{y}_i^{\mathscr{V}} - \boldsymbol{y}_j^{\mathscr{V}}||_2$ must also be below a user-specified threshold $y_{\max}$.

5.2 RSO Custody

To maintain RSO custody, the MAP module computes an attitude profile to align the camera boresight $\hat{\boldsymbol{z}}^{\mathscr{V}}$ with the predicted relative position vector $\delta\boldsymbol{r}^{\mathscr{V}}$ of the target at the next scheduled image epoch t_{im}. The target is autonomously selected from a list $\boldsymbol{T}$.

For each target i in the list, the ISD $||\delta\boldsymbol{r}^{\mathscr{V}}||_2$ is computed. If i is predicted to be in eclipse or its ISD is above a threshold $d_{\max}$ (making it too dim to be observed), the target is invalidated. Otherwise, the rotation ${}^{\mathscr{I}}\overrightarrow{\boldsymbol{R}}_i^{\mathscr{V}}$ which fulfills ${}^{\mathscr{I}}\overrightarrow{\boldsymbol{R}}_i^{\mathscr{V}} \delta\boldsymbol{r}_i^{\mathscr{I}} = [0, 0, 1]^\top$ is computed. If this rotation causes the camera to be blinded by the Earth, Sun or Moon, the target is also invalidated.

Each Starling CubeSat carries two star trackers, pointed in approximately the velocity and anti-velocity directions. If the target is trailing the observer (i.e. $\delta r_y^{\mathscr{B}} \leq 0$, the backwards-facing tracker is chosen, and vice-versa for leading targets. The body-frame Euler angle vector $\boldsymbol{e}_i^{\mathscr{B}} = [\phi_i, \theta_i, \psi_i]$ which aligns the chosen tracker with the target is then computed, along with the control effort required to achieve $\boldsymbol{e}_i^{\mathscr{B}}$ with respect to the current attitude. A target score is computed using the ISD and control effort values.

If one or more valid targets exist in the list, the target which minimizes the scoring criteria is chosen. Otherwise, the satellite remains in its default attitude. Between 1 and 50 potential targets were provided in flight and in practice, minimizing target distance (or maximizing brightness) was preferred at the expense of larger control effort.

5.3 RSO State Refinement

State refinement for the observer and RSOs is performed in SOD using an unscented Kalman filter (UKF) framework. In comparison to an extended Kalman filter, the UKF is able to incorporate fully nonlinear dynamics and measurement models and can more accurately approximate higher-order behavior. The resulting observability improvements are beneficial for robust angles-only state convergence [7]. In flight, observer and RSO states were handled independently in the filter to minimize state vector size and corresponding computation time, at the cost of losing cross-covariance information. The orbit estimate of the observer was updated using PVT solutions from GPS and the numerical propagator of Section 2. Flight experiments were also performed in which the observer's orbit was refined using only optical measurements to identified RSOs [9], but not at the same time as RSO refinement. RSO states were instead updated using bearing angle observations provided by IMP.

When an RSO is observed for the first time on orbit, its state in the UKF must be initialized. Its TLEs in the onboard catalog are propagated to t_{im} using SGP4, and its uncertainty is initialized using a time-dependent empirical model similar to Equation 6. If the RSO has been observed previously (i.e. a state estimate exists), a UKF time update is performed. Its filter state and covariance are propagated to t_{im} by applying SGP4 through an unscented transform, and process noise is incorporated through a time-dependent orbit element noise model.

The typical UKF measurement update [40] is then completed. Modeled bearing angles are computed for each sigma point as per Section 2, and are applied in combination with the measured bearing angle to refine the RSO's estimated state and covariance. In flight, a constant bearing angle measurement noise was assumed for RSOs, of approximately $40''(1\sigma)$ in azimuth and elevation. The refined RSO states are subsequently stored in the onboard catalog in the form of TLEs and associated covariance matrices. During flight experiments only the $[n, e, i, \Omega, \omega, M]$ TLE components were updated online; the drag term and rates of change of mean motion were assumed constant over the experiment period.

5.4 RSO State Initialization

In contrast to the single-observer BOD algorithm described in previous StarFOX work [31], the new BOD algorithm leverages bearing angles shared between multiple observers to improve observability and speed of convergence. Consider the model $\boldsymbol{y}_i^{\mathscr{I}}(\boldsymbol{x}, t)$ that provides bearing angles of target i as a function of system state $\boldsymbol{x}$ at time t. Bearing angles may be equivalently expressed as line-of-sight (LOS) unit vectors $\hat{\boldsymbol{v}}_i^{\mathscr{I}}(\boldsymbol{x}, t)$ from observer to target,

$$\hat{\boldsymbol{v}}_i^{\mathscr{I}} = [v_x^{\mathscr{I}}, v_y^{\mathscr{I}}, v_z^{\mathscr{I}}]^\top = \frac{\boldsymbol{v}_i^{\mathscr{I}}}{||\boldsymbol{v}_i^{\mathscr{I}}||_2} \tag{9}$$

where $\hat{\cdot}$ indicates a unit vector quantity. Era-Core operates on batches of LOS vectors provided across N_y epochs $t_1, ..., t_{N_y}$, collectively denoted $\boldsymbol{t}$. A measurement batch $\boldsymbol{z}_i^j(\boldsymbol{x}, \boldsymbol{t})^{\mathscr{I}} \in \mathbb{R}^{2N_y}$ is

$$\boldsymbol{z}_i^j(\boldsymbol{x}, \boldsymbol{t})^{\mathscr{I}} = \begin{pmatrix} \hat{\boldsymbol{v}}_i^{\mathscr{I}}(\boldsymbol{x}, t_1) \\ \vdots \\ \hat{\boldsymbol{v}}_i^{\mathscr{I}}(\boldsymbol{x}, t_{N_y}) \end{pmatrix} \tag{10}$$

where subscript $\cdot_i$ indicates target object i and superscript $\cdot^j$ indicates observer j. Assume j produces a batch $\boldsymbol{z}_i^j(\boldsymbol{x}, \boldsymbol{t})$ of local measurements of i, and receives a batch $\boldsymbol{z}_i^k(\boldsymbol{x}, \boldsymbol{t})$ of measurements from observer k over the ISL. The observer may express each LOS vector in the RTN frame of the measurement's source, producing $\boldsymbol{z}_i^j(\boldsymbol{x}, \boldsymbol{t})^{\mathscr{R}^j}$ and $\boldsymbol{z}_i^k(\boldsymbol{x}, \boldsymbol{t})^{\mathscr{R}^k}$.

The target's relative position vector $\delta \boldsymbol{r}_i^{\mathscr{R}}$ can be normalized by the observer's orbit radius to form the non-dimensional vector $\delta \bar{\boldsymbol{r}}_i^{\mathscr{R}} = [\delta \bar{x}, \delta \bar{y}, \delta \bar{z}]^\top$, where $\bar{\cdot}$ indicates a non-dimensional quantity. Each LOS vector $\hat{\boldsymbol{v}}_i^{\mathscr{R}}$ is parallel with $\delta \bar{\boldsymbol{r}}_i^{\mathscr{R}}$ at its measurement epoch. A set of constraint equations is obtained from their cross product, with

$$\hat{\boldsymbol{v}}_i^{\mathscr{R}} \times \delta \bar{\boldsymbol{r}}_i^{\mathscr{R}} = \begin{bmatrix} v_z \delta \bar{y} - v_y \delta \bar{z} \\ v_x \delta \bar{z} - v_z \delta \bar{x} \\ v_y \delta \bar{x} - v_x \delta \bar{y} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}. \tag{11}$$

Furthermore, a target's relative state with respect to an observer is characterized by a set of six parameters, here the quasi-nonsingular ROE. The linear mapping between relative position $\delta \boldsymbol{r}^{\mathscr{R}}$ and corresponding ROE $\delta \boldsymbol{\alpha}$ has the form

$$\bar{x}^{\mathscr{R}} = \boldsymbol{b}_x^\top(\theta) \delta \boldsymbol{\alpha} \qquad \bar{y}^{\mathscr{R}} = \boldsymbol{b}_y^\top(\theta) \delta \boldsymbol{\alpha} \qquad \bar{z}^{\mathscr{R}} = \boldsymbol{b}_z^\top(\theta) \delta \boldsymbol{\alpha}. \tag{12}$$

The vectors $\boldsymbol{b}_{x,y,z} \in \mathbb{R}^6$ are functions of the observer's true argument of latitude θ and capture the time variation of the relative position coordinates.

Following the approach of Willis and D'Amico [41], a linear system is formed by substituting Equation 12 into Equation 11. Each LOS vector provides two linearly independent constraints; thus, three LOS vectors are needed to fully determine the six-dimensional relative state vector. The resulting system for three LOS vectors at $\boldsymbol{t} = [t_1, t_2, t_3]$ from observer j is compactly written as $V^j \delta \boldsymbol{\alpha}_i^j = \boldsymbol{0}$, where $\delta \boldsymbol{\alpha}_i^j$ denotes the ROE of i with respect to j.

Consider a second observer k with $V^k \delta \boldsymbol{\alpha}_i^k = \boldsymbol{0}$. Since $\delta \boldsymbol{\alpha}_i^k = \delta \boldsymbol{\alpha}_k^j - \delta \boldsymbol{\alpha}_i^j$, this may be rearranged as $V^k \delta \boldsymbol{\alpha}_i^j = V^k \delta \boldsymbol{\alpha}_k^j$. The two systems can be combined to solve for a unique target ROE vector $\delta \boldsymbol{\alpha}_i^j$, when given the LOS batch from j, the LOS batch from k, an absolute orbit estimate for j, and an absolute orbit estimate for k.

This framework is combined with a sampling approach to produce an accompanying uncertainty estimate, and to perform averaging for the effects of measurement noise on the solution. Consider randomly selecting N_b measurements from batch $\boldsymbol{z}_i^j(\boldsymbol{t}^j)$ and N_b measurements from batch $\boldsymbol{z}_i^k(\boldsymbol{t}^k)$, where $N_b \geq 3$. The resulting linear system can be solved to produce a solution sample $\delta \boldsymbol{\alpha}_i^j(1)$. This process is repeated to generate N_s state samples $[\delta \boldsymbol{\alpha}_i^j(1), \delta \boldsymbol{\alpha}_i^j(2), \ldots, \delta \boldsymbol{\alpha}_i^j(N_s)]$ and the mean and variance of the sample set provides the output ROE mean and variance for target i.

6. FLIGHT RESULTS

6.1 RSO Identification

During flight experiments from April 2025 to July 2026, 16003 RSO detections were made in star tracker images on SV2 and SV4. 3623 unique RSOs were identified. The farthest object identified was NORAD ID 16910, a Japanese H-1 rocket booster with a radar cross section of 29m^2 and ISD of 3948 km, on 9/28/25. The closest non-Starling object identified was NORAD ID 55456, a Starlink satellite with a radar cross section of 14 m^2 and an ISD of 125 km, on 6/13/25. The most frequently observed non-Starling objects were NORAD ID 56323, a Starlink satellite, which was identified in 75 images, and NORAD ID 39260, a Fengyun 3C satellite, which was identified in 49 images. Stored 20×20 pixel regions of interest for these objects are provided in Figure 7. Multiple detections of the same object are overlaid, and pixel brightness has been scaled for visibility.

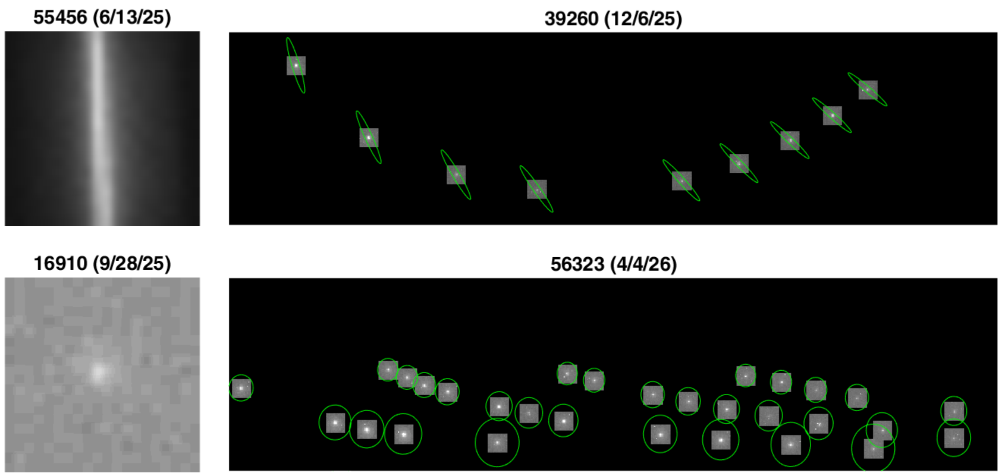

Fig. 7: Star tracker regions of interest for four identified RSOs. Green ellipses display the 3σ uncertainty region used for RSO identification.

The correctness of identifications is assessed by computing $e_{\text{ang}} = ||\boldsymbol{y}^{\mathscr{V}}_{\text{meas}} - \boldsymbol{y}^{\mathscr{V}}_{\text{ref}}||_2$, i.e. the residual between the measured bearing angle of the RSO and its bearing angle as determined by a reference truth. The angular residual can be scaled by the ISD to produce a geometric error in meters, $e_{\text{geo}} = e_{\text{ang}}||\delta \boldsymbol{r}_{\text{ref}}||_2$. A catalog residual is also defined to express the difference between the onboard catalog position and reference position, with $e_{\text{cat}} = ||\boldsymbol{r}^{\mathscr{I}}_{\text{cat}} - \boldsymbol{r}^{\mathscr{I}}_{\text{ref}}||_2$.

The source of reference truth depends on the object type. For the Starling swarm, the reference truth is produced via batch post-processing of onboard GPS data, with expected position errors of less than 10 m. For Starlink satellites, the reference truth is obtained from predicted ephemeris published by SpaceX, with errors on the order of tens to hundreds of meters. For other RSOs, the reference truth is obtained from the Space-Track TLE solution closest to t_{im}, with errors on the order of hundreds of meters to kilometers. An identification is considered correct if $e_{\text{ang}} \leq 0.001$, or incorrect if $e_{\text{ang}} \geq 0.01$. All other identifications are considered ambiguous. The measurement noise floor of the Starling star tracker is approximately 20-40" (1σ) which corresponds to $e_{\text{ang}} = 0.0001$.

Table 4 presents statistics for on-orbit identifications from April 2025 to June 2025. The overall correct match rate was 81%, being lowest for Starlink RSOs and highest for 'other' RSOs. The 'other' category possesses the largest ISD on average because it contains larger satellites visible at farther distances. Thus, their uncertainty regions are smallest in the FOV, with a smaller chance of being assigned to an unrelated nearby object. The Starling category possesses the lowest ISD, but their relative velocities are low, meaning the size of the covariance ellipse remains bounded. The Starlink category often features large components of relative velocity perpendicular to the camera boresight, producing large, elongated uncertainty regions. Thus, there is a larger probability of the ellipse encompassing an unrelated object. In practice, the proportion of incorrect matches can be reduced by lowering the $\varepsilon_{\text{assign}}$ and y_{max} thresholds.

Table 4 also presents mean angular, geometric and catalog residuals. Angular residuals were lowest for Starling, which suggests its TLE solutions are more accurate than the rest of the catalog, on average. Starlink angular residuals are somewhat larger due to more frequent misidentifications. Catalog residuals are also largest for Starlink, which implies

Table 4: RSO identification and measurement statistics.

	Correct Match	Ambiguous Match	Incorrect Match	Correct Match Rate	e_{ang}	e_{geo}	e_{cat}
					(Mean $\pm$ 1σ, correct matches only)		
Starling	1744	209	188	81.5%	$30'' \pm 27''$	14 ± 14 m	439 ± 324 m
Starlink	202	82	23	65.8%	$80'' \pm 54''$	62 ± 24 m	2844 ± 2024 m
Other	599	69	16	87.6%	$60'' \pm 44''$	64 ± 25 m	1015 ± 943 m
All	2545	360	227	81.3%	$41'' \pm 39''$	17 ± 20 m	532 ± 653 m

significant differences in fidelity between Starlink TLE solutions, as propagated onboard, and Starlink ephemeris data, used for reference truth on the ground. Figure 8 presents histograms of the geometric residuals across all identifications. Similar patterns are observed, in that Starling residuals are lowest and Starlink residuals are highest.

The truest picture of measurement accuracy is provided by the Starling measurements because their reference truth is the most accurate among those surveyed. In Table 4, the mean e_{geo} is approximately 14 m per 100 km of ISD. This approaches GPS PVT accuracy for nearby targets. Furthermore, the majority of non-Starling measurements in Figure 8 displayed errors less than 100 m per 100 km of ISD. On-orbit measurements are therefore valuable for enhancing SSA accuracy, given that position uncertainties in TLE products are typically much larger [25].

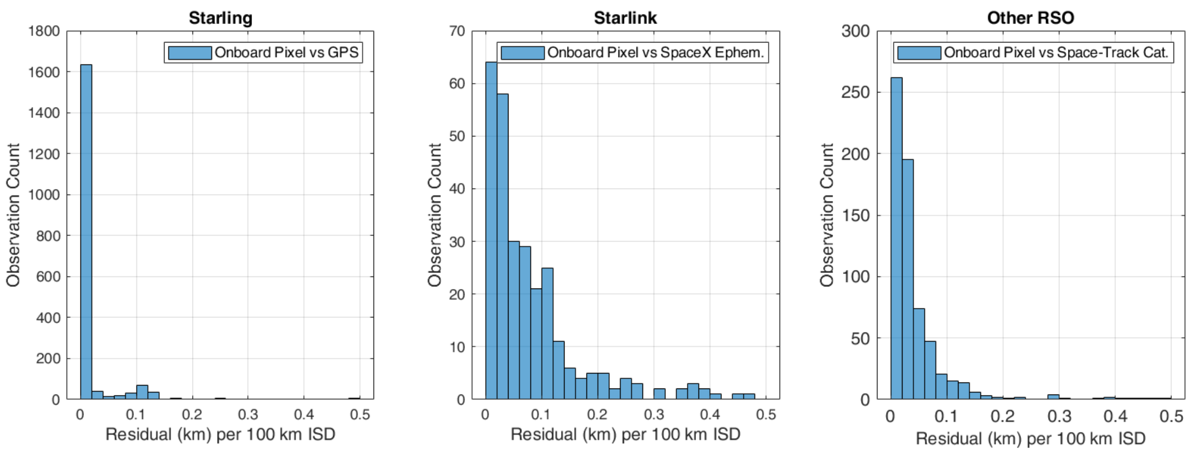

Fig. 8: Geometric residuals between the measured bearing angle and the object's reference position.

6.2 RSO Custody

Maintaining persistent visual custody of a single spacecraft was explored for SV4 tracking SV1. As per Table 3, SV1 possesses a large relative orbit with respect to SV4, with periodic in-plane oscillations of ~16 km and out-of-plane oscillations of ~63 km. For the typical Starling ISD of 50-150 km between swarm members, this corresponds to bearing angle magnitudes of 20-50° when measured with respect to the velocity direction $\pm\hat{\boldsymbol{y}}^{\mathscr{W}}$. Thus, if the star tracker boresight was kept in its default velocity-aligned attitude, SV1 was generally outside the FOV.

During an experiment on 9/20/25, 829 star tracker images were taken over a 19-hour period. The target was expected to be visible for a maximum of twelve minutes per orbit in the default imaging attitude, which proportionally suggests approximately 103/829 images with target presence, discounting other optical constraints. In practice, SV1 was observed 239 times, displaying a 2.3x increase in measurement availability. Figure 9 plots bearing angles of SV1 with and without active target tracking, as well as the commanded Euler angles for the SV4 body frame. Gaps were present between periodic measurement arcs due to visibility issues (such as eclipse and blinding, affecting 420 images), large ISD, and other spacecraft tasks such as ground contacts that disrupted observations. However, measured bearing angles were consistently within 0.5 degrees of the image center, demonstrating accurate pointing and custody.

Subsequent experiments on 1/17/26 and 1/18/26 explored autonomous custody of an RSO list. On each day, the list provided to the satellite was made up of objects predicted to have a low relative velocity, to provide the most opportunities for repeated detections. The observer was then allowed to independently select the target from the list

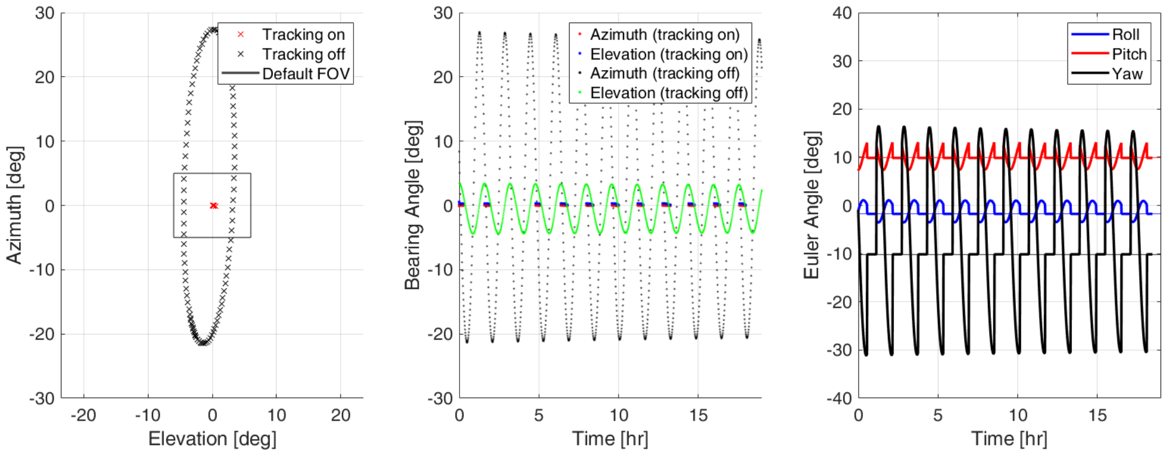

Fig. 9: Target bearing angles in the image plane (left, center) and commanded body-frame Euler angles (right).

which minimized the ISD. In total, 1949 observations were planned onboard encompassing 75 unique RSO targets. Approximately 1030 RSO observations were achieved, in that the desired RSO was identified in the image. Figure 10 presents two examples of attitude profiles used to image RSOs. The first, NORAD ID 57959, was imaged 9 times with a relative velocity of 2.3 km/s. The second, NORAD ID 37789, was imaged 45 times with a relative velocity of 0.3 km/s. The most common reasons for missing observations included poor RSO visibility of faraway targets and disruption of tracking by charging orbits. The variety of target RSO orbits meant that Euler angles varied frequently and significantly, resulting in poor conditions for solar power conversion. One in every three orbits was therefore reserved for charging with a fixed attitude profile, to maintain safe battery levels throughout the experiment.

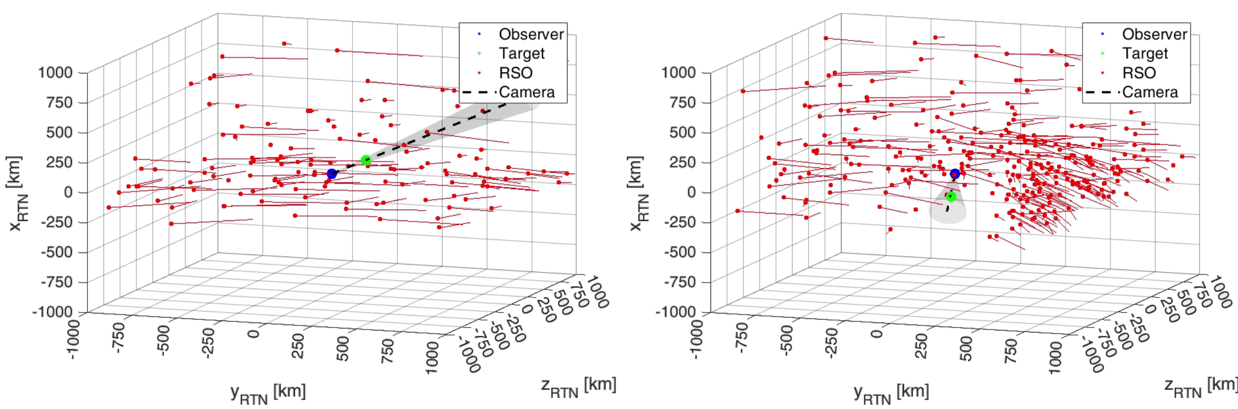

Fig. 10: Visualization of observation geometry and the RSO neighborhood for two on-orbit images.

6.3 RSO Refinement

Refinement of RSO catalog orbits was demonstrated during flight experiments on 6/13/26 and 6/15/26. Starlink satellites were used as targets for refinement because vehicle ephemerides (including state uncertainties) are publicly available from SpaceX, providing reference orbit data with more accuracy and consistency than standard TLEs. Four data sources were leveraged to assess the quality of RSO orbit information refined onboard. First are the TLEs provided in the onboard catalog, denoted 'pre', which are typically 2-7 days old before being used on board. Second are TLEs downloaded from Space-Track on the day of the experiment, denoted 'new', as an example of less stale orbit data. Third are Starlink ephemerides downloaded from SpaceX on the day of the experiment, denoted 'ref', as a reference for the target's true orbit. Fourth are the TLEs refined on board, denoted 'post'.

Consider the position error in onboard TLEs post-refinement $e_{\text{post}} = ||\boldsymbol{r}_{\text{post}} - \boldsymbol{r}_{\text{ref}}||_2$, the position error in the newer on-ground TLEs $e_{\text{new}} = ||\boldsymbol{r}_{\text{new}} - \boldsymbol{r}_{\text{ref}}||_2$, and the position error in the stale onboard catalog $e_{\text{pre}} = ||\boldsymbol{r}_{\text{pre}} - \boldsymbol{r}_{\text{ref}}||_2$. Figure

11 plots histograms of the error differences $e_{\text{post}} - e_{\text{new}}$ and $e_{\text{post}} - e_{\text{pre}}$ (such that a negative error delta indicates an improvement in position knowledge). Statistics regarding these error metrics and filter residuals are given in Table 5.

On-board refinement led to significant improvements to position error in most cases, of up to 25 km in Figure 11. Mean position errors were lowest for e_{post}, improving upon e_{new} by a factor of 2.3x and improving upon e_{pre} by a factor of 4.0x. Standard deviations of error were also lowest when refinement is performed onboard. Thus, online refinement is able to significantly enhance the quality, consistency, and timeliness of RSO orbit knowledge. Filter residuals also showed large decreases after onboard measurement updates.

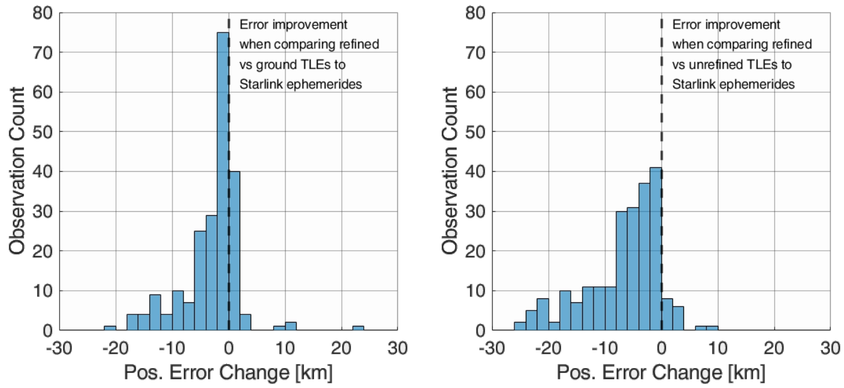

Fig. 11: Changes in RSO position error when comparing onboard refinement to more current on-ground TLE solutions (left) and the unrefined onboard catalog (right).

Table 5: RSO refinement error statistics.

	Mean	St. Dev. (1σ)	Median	Minimum	Maximum
Position Error e_{post}	2.29 km	5.51 km	1.00 km	0.12 km	60.05 km
Position Error e_{new}	5.33 km	6.62 km	2.80 km	0.08 km	42.74 km
Position Error e_{pre}	9.11 km	7.55 km	7.04 km	0.08 km	51.76 km
Pre-fit Residuals	0.346°	0.369°	0.198 °	0.0010°	1.479°
Post-fit Residuals	0.012°	0.015°	0.008 °	0.0003°	0.152°

In some cases, onboard refinement appeared to increase orbit state errors, likely due to misidentification of RSOs. The low image sample rate meant RSOs often had to be identified using a single image, and the age of the onboard catalog created large state uncertainties as time progressed. Figure 12 presents a refinement example featuring a large bearing angle uncertainty prior, such that the semimajor axis of the uncertainty ellipse spans multiple degrees. Relaxing identification thresholds (as outlined in Section 5.1) to increase the number of potential observations led to a corresponding increase in ambiguous RSO matches which could not be robustly detected by the filter. It is also possible that Starlink maneuvers disrupted some ephemerides, adding to perceived errors during analysis; for example, if $t_{\text{maneuver}} < t_{\text{post}} < t_{\text{ref}}$ or $t_{\text{post}} < t_{\text{maneuver}} < t_{\text{ref}}$, there is a discrepancy in whether the refined or reference orbit data reflects the maneuver effect.

6.4 RSO Initialization

Fast initialization of a new RSO orbit estimate was achieved by SV2 and SV4 on 12/6/25. First, observers independently collected optical measurement batches of SV1. Every 10 minutes, measurement histories were shared between observers over the ISL. Every 90 minutes, the BOD algorithm of Section 5.4 was executed and if sufficient measurements were present, an output state estimate was generated. Due to the various errors sources present in the system (e.g. observer orbit uncertainties and target measurement uncertainties), approximately 20 minutes of measurements were necessary to generate a robust target orbit estimate. The process is visualized in Figure 13.

Figure 14 plots the 50 output state samples used to generate the mean and uncertainty (1σ) for SV1's relative orbit estimate, together with the true state. The true ROEs lie within the 1σ ellipse around the estimated ROEs, indicating good state convergence. When transformed into RTN position-velocity, the output position error was 1.72 km with a

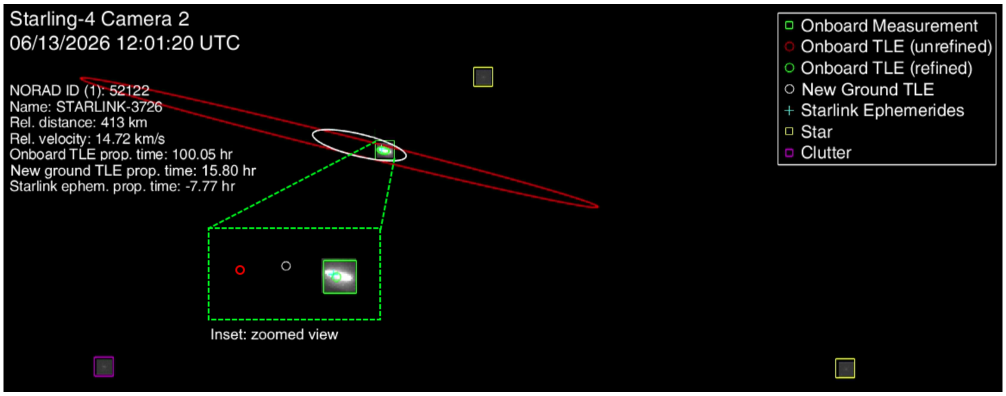

Fig. 12: A reconstructed star tracker image with uncertainty regions pictured for refinement of STARLINK-3726.

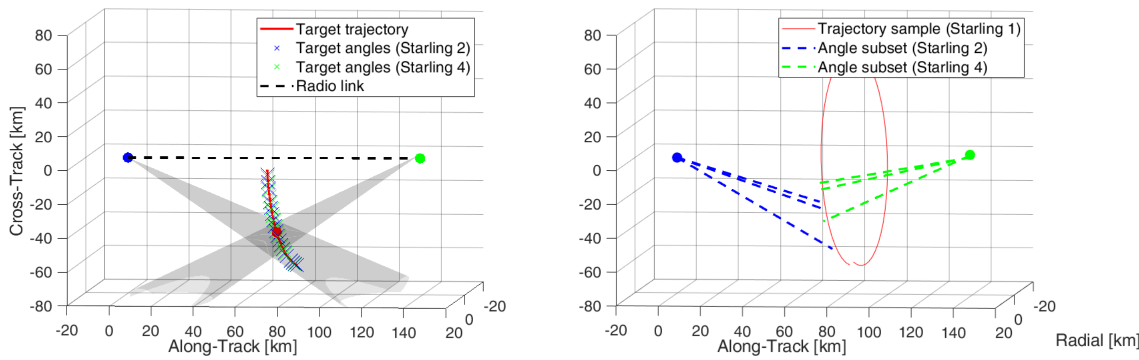

Fig. 13: The captured bearing angle batches for SV1 orbit initialization (left) and a single output state sample from the BOD algorithm (right).

1σ uncertainty of 1.48 km, and the output velocity error was 0.712 m/s with a 1σ uncertainty of 1.55 m/s. Position errors were approximately 2% of the 68 km ISD and errors remained well within 3σ state uncertainties.

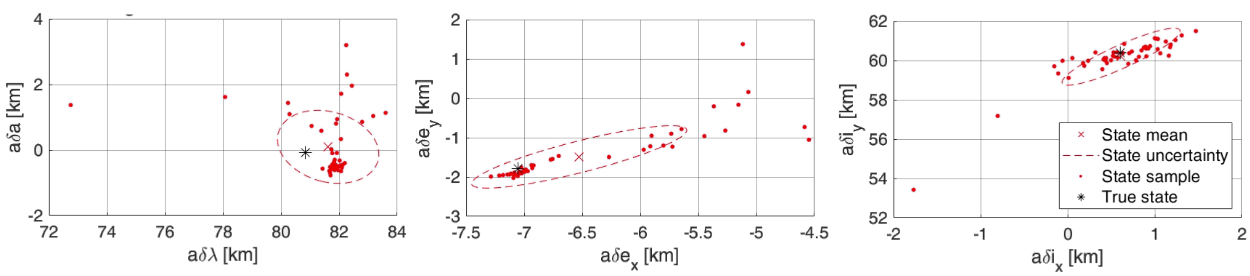

Fig. 14: SV1 ROE state samples and output mean and covariance generated by the multi-observer BOD algorithm.

The single-observer BOD algorithm previously published by the authors required two or more hours of measurements to achieve convergence in flight [31]. In comparison, the multi-observer system reduces the required observation time by a factor of >6, demonstrating the power of a multi-sensor architecture. Observation time could likely be reduced further if higher-quality measurements were available; bearing angle noise is strongly detrimental to estimation performance when using short measurement arcs [42], which for Starling is on the order of $40''$.

Upcoming flight experiments will demonstrate additional features of the Era-Node system. Starling experiments in 2026-27 will explore distributed multi-observer catalog refinement over the ISL, as well as online conjunction assessment and collision avoidance maneuver planning with respect to the RSO catalog.

7. CONSTELLATION SIMULATION

Feasibility of the Era-Cloud component of the architecture is verified through high-fidelity simulations of the complete space-based SSA network. Proliferated observers are distributed over LEO altitudes between 350-2000 km in near-circular orbits, drawn from existing RSO orbits in the Space-Track catalog. This semi-random choice of orbits emulates the proposed virtual constellation, as opposed to a specifically-designed SSA constellation.

Four case studies are defined to characterize performance. Case 1 features 140 observers using star trackers (with the parameters of the NST in Table 1) to enable onboard autonomy. Case 2 applies a custom 'Era-Node' sensor with a FOV of $40^\circ \times 30^\circ$and a maximum detectable magnitude of 8.0, which can be conceptualized as a mosaic of multiple star trackers. Cases 3 and 4 scale up to 1090 observers for the star tracker and Era-Node respectively. Each observer aligns its sensor with its velocity direction with a sample rate of 0.2 Hz.

7.1 Simulation Environment

Experiments apply EraDrive's 'Era-Twin' simulation environment. The Space-Track database [24] is used to provide TLEs for the RSO population, and the ESA DISCOS database [43] is used to provide RSO area data. To generate reference orbit data, all TLEs are propagated to measurement epochs using SGP4. Observer orbits are propagated using numerical integration of the GVEs, with perturbations including a 10x10 GGM01S spherical harmonic gravity model and Harris-Priester atmospheric density with cannonball drag. Subsequently, the bearing angles for each RSO as seen from each observer are generated, as well as a visual magnitude estimate. The magnitude computation incorporates the observer-target-sun phase angle and flux contributions from multiple sources. If the bearing angle lies within the sensor's FOV and estimated magnitude is brighter than its detection limit, a noisy measurement is generated according to the sensor's noise model. All measurements of visible RSOs across the simulation period are stored to enable computation of performance metrics. High-level data flow is diagrammatically represented in Figure 15.

A prototype centralized estimation algorithm was also implemented to provide an indication of the achievable estimation accuracy. It is an extension of the BOD algorithm used during the StarFOX experiment, which applied iterative batch least squares refinement to generate orbit estimates using angles-only measurement batches in flight [31]. In comparison to that work, the extended algorithm incorporates bearing angles from multiple observers during state refinement, which greatly enhances observability and eliminates the need for target range sampling.

For each RSO, an initial guess for its state is obtained by corrupting its original catalog TLEs with Gaussian errors. All stored bearing angles and accompanying observer orbit estimates, as would be downlinked from the virtual constellation, are then processed to compute an orbit estimate. An accompanying covariance estimate is generated through numerical methods [12] with uncertainty contributions from three primary sources. First, it is assumed observers have access to GNSS such that a geometric uncertainty 20 m (1σ) is present in each of their orbit elements. An angular uncertainty of 20" (1σ) is assumed in each bearing angle. Finally, dynamics uncertainty is incorporated by including unmodeled RTN accelerations $\boldsymbol{d}^{\mathscr{R}} = [1,1,1]^\top \times 10^{-7}$ m/s^2 which act on the RSO state. Orbit propagation within the batch estimator is handled using an analytic model which incorporates the J2 gravity term [17]. The output of this process is considered an upper bound on the achievable error due to the relative simplicity of the prototype algorithm, the low fidelity of the dynamics model, and the conservative inclusion of error sources.

7.2 Simulation Results

Simulations encompass a two-hour period beginning on June 30, 2026 00:00:00 UTC. The RSO catalog consists of 30653 objects and of these, 26676 are considered LEO objects with altitudes of 2000 km or lower. Each observer generates 1440 images with which to cover the LEO environment. Figure 16 presents a visualization of observer and target locations in LEO at the start of the Case 4 simulation, along with the simulated view for one observer satellite.

Statistics for constellation-wide SSA performance are presented in Table 6. Coverage is defined as the percentage of all objects tracked in a category, where 'large' refers to a cross-sectional area greater than 0.5 m^2. The median revisit rate (MRR) is the time between measurements of the same object, and the mean maximum time unobserved (MMTU) refers to the longest gap between successive measurements. Accuracy refers to the error in the output estimate from the prototype estimation algorithm, compared to the simulation truth. Uncertainty is the output covariance (3σ) produced by the algorithm. Initial state errors for BOD were on the order of 5 km in position and 5 m/s in velocity (1σ).

The fully proliferated Case 4 scenario achieves 95% coverage of the LEO environment in only two hours, with 100% coverage of large LEO objects. The MRR is less than two minutes, indicating persistent RSO custody and the potential

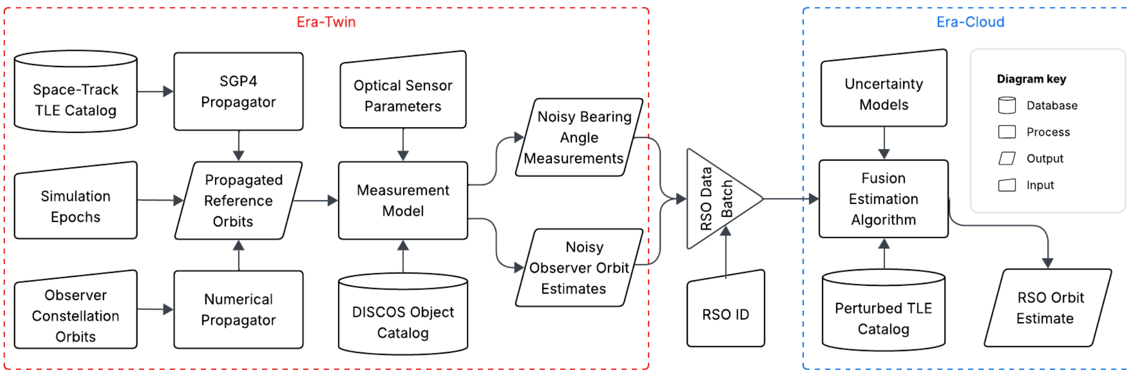

Fig. 15: High-level data flow for the Era-Twin SSA simulation.

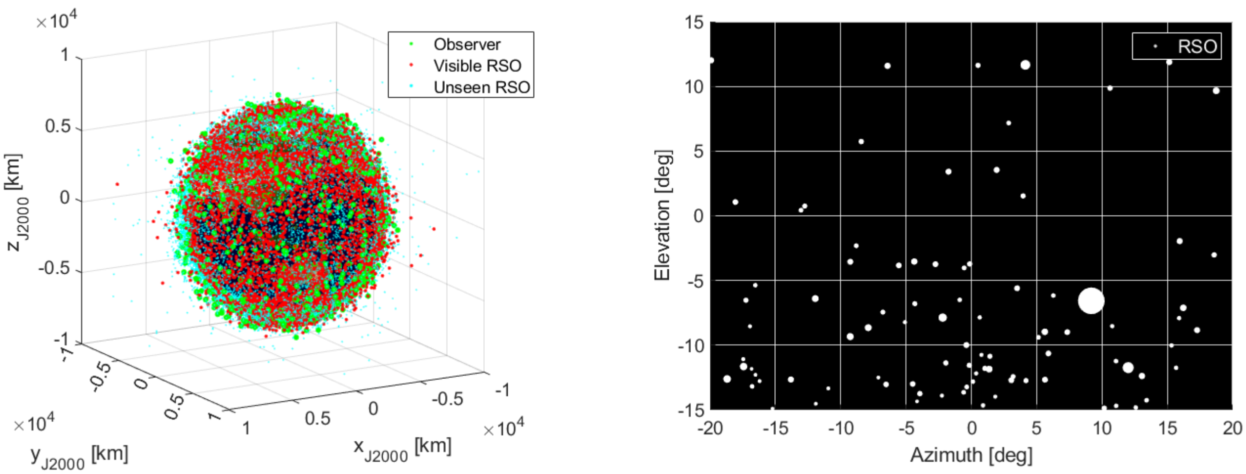

Fig. 16: The LEO RSO neighborhood for Case 4 (left) and simulated FOV for one observer at 600 km altitude (right).

to responsively track RSOs performing unknown maneuvers. The position and velocity accuracies achieved are on the order of tens of meters, as are the uncertainties produced by the prototype algorithm. This level of orbit knowledge, achieved in a two-hour period with entirely passive measurements, improves significantly over TLE products today.

Case 2 also demonstrates strong performance, despite an 87% reduction in constellation size. The proposed Era-Node sensor provides a sufficiently large detection region (while retaining low size, weight and power properties) to enable high coverage of large LEO objects, a relatively frequent MRR, and estimation accuracies similar to Case 4. Thus, useful SSA outcomes can be achieved while scaling towards complete proliferation of the concept.

It is also useful to examine performance when leveraging a standard star tracker. No additional hardware is needed, and SSA capabilities could even be enabled through over-the-air software updates of existing satellites, which greatly eases the cost of proliferation (though the number of available measurements is limited in comparison). Nevertheless, Case 1 provides 188000 measurements of over 9000 objects in two hours, a non-trivial data volume which could provide strong benefits when integrated with other sources in a broader SSA architecture. Case 3 provides much stronger standalone performance: 96% coverage of large LEO objects, an MRR of less than 4 minutes, and estimation accuracies comparable to Case 4. The decreased measurement availability does lead to larger uncertainties however, and multiple star trackers are needed to provide the same coverage as a single Era-Node sensor.

Overall, simulations indicate that proliferated, decentralized space-based SSA is a viable alternative or additive to traditional architectures. Existing satellite platforms with star trackers, new platforms carrying optical sensors for autonomy, and ground-based infrastructure can be combined to provide timely global SSA coverage and highly accurate

Table 6: SSA simulation performance metrics.

	Case 1	Case 2	Case 3	Case 4
Observers	140	140	1090	1090
Sensor	Star tracker	Era-Node	Star tracker	Era-Node
Objects Tracked	9072	18516	18167	25665
Coverage (all LEO objects)	33.69%	68.48%	67.50%	94.88%
Coverage (large LEO objects)	55.16%	97.91%	95.53%	99.99%
Minimum Size Observed (m^2)	0.015	0.015	0.015	0.0058
Measurements Taken	188612	4080781	1835479	37150360
Median Revisit Rate (min)	35.38	7.58	3.75	1.67
Mean Max. Time Unobserved (min)	75.16	50.31	50.79	39.80
Median Position Accuracy (m)	100.13	65.02	68.79	61.39
Median Position Uncertainty (m)	946.44	160.36	268.24	89.59
Mean Velocity Accuracy (mm/s)	73.42	64.03	63.95	61.81
Mean Velocity Uncertainty (mm/s)	844.30	157.49	230.18	89.99

RSO orbit knowledge. Promising outcomes can be achieved with relatively few vehicles and performance increases as the system scales with the RSO population.

Some key questions remain, such as the aggregated latency between observer data downlink, on-ground generation of SSA catalogs, and subsequent dissemination, and how this compares with existing architectures. Future simulations will be extended to include modeling and assessment of data latency, as well as analysis of SSA in GEO and characterization of how space-based sensors can improve tracking of the small debris object population.

8. CONCLUSION

EraDrive's novel architecture for autonomous space situational awareness (SSA) and space traffic coordination (STC) consists of two complementary processes: a decentralized online procedure which leverages a virtual constellation of proliferated space-based observers to gather SSA data, and a centralized offline procedure which operates on the ground to generate global SSA catalogs. Each observer carries an 'Era-Node' module featuring a processor, a crosslink and a camera, which may be a star tracker or a more capable sensor. Onboard software is used to opportunistically detect and identify resident space objects (RSOs) in images, gather bearing angle measurements, refine local RSO catalogs, and generate maneuver plans for orbit maintenance and collision avoidance. Crosslinks are used to share data with other Era-Nodes and the ground. In doing so, satellites are provided with responsive awareness of their environment and can coordinate downstream maneuver planning and collision avoidance with improved autonomy and safety. The offline 'Era-Cloud' layer consists of a centralized database which receives measurement and orbit data from all Era-Nodes (and potentially other external sources) to generate updated global RSO catalogs. The incorporation of space-based measurements enables more accurate and persistent tracking of the RSO population.

Feasibility of the Era-Node layer is verified through flight demonstrations of autonomous SSA aboard the NASA Starling swarm. Real-time SSA capabilities were exhibited by EraDrive's software operating on Starling hardware, including matching of transient RSO targets to identities in the space object catalog; generation and execution of attitude plans to center targets in the field of view and enable persistent custody; refinement of the onboard RSO catalog using new measurements, with demonstrable improvements in orbit knowledge; and initialization of new orbit estimates for RSOs using bearing angles shared between multiple observers. Flight experiments confirm that autonomous space-based SSA can be realized using low size, weight, power and cost hardware, with concrete operational benefits.

Feasibility of the Era-Cloud layer is verified through high-fidelity simulations of the complete RSO population as tracked by the virtual Era-Node constellation. Simulations featuring up to 1090 Era-Nodes achieve almost 100% coverage of low Earth orbit in two hours, with revisit rates on the order of minutes and orbit determination accuracies of tens of meters. The resulting SSA products display better timeliness and quality than those typically available today. The system also displays promising properties with regards to scaling, and relatively few space-based observers can provide a large volume of additional measurements to existing SSA architectures. In comparison to current ground-based sensing or ground-in-the-loop paradigms, the proposed system provides multilayered benefits to the safety and autonomy of individual satellites as well as space sustainability at scale.

9. ACKNOWLEDGMENTS

This work was supported by NASA contract number 80ARC025CA006 titled "Autonomous Space-Based Situational Awareness through EraDrive Autonomy". The authors thank Soon S. Hwang and the entire NASA Ames Research Center Starling team for their tireless support of StarFOX follow-on efforts.

10. REFERENCES

[1] Andrew Falle, Ewan Wright, Aaron Boley, and Michael Byers. One million (paper) satellites. *Science*, 382(6667):150–152, 2023.

[2] Matthew Jenkins. Maneuvering Towards Safe and Responsible Rendezvous and Proximity Operations in Space. *Space Education & Strategic Applications*, 4(4), 2024.

[3] Simone D'Amico, Tycho Bogdanowitsch, and Rebecca Wang. Space Safety and Sustainability, Part I: Past and Present Practices. Policy report, Hoover Institution, Stanford University, September 2025. Hoover Technology Policy Accelerator.

[4] Daniel O'Connell, Jean-Claude Leclerc, Noemi Giammichele, Ryan Comeau, Narendra Gollu, Naron Phou, and Frederic Pelletier. Concept of Operation and Initial Performance Summary of the NorthStar Space-Based Optical SSA System. In *Proceedings of the 25th Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS)*, 2024.

[5] SpaceX. Stargaze: SpaceX's Space Situational Awareness System, January 2026.

[6] D. A. Vallado and W. D. McClain. *Fundamentals of Astrodynamics and Applications*. Microcosm Press, Hawthorne, California, 4 edition, 2013.

[7] 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.

[8] S. D'Amico, J.-S. Ardaens, G. Gaias, et al. Noncooperative Rendezvous Using Angles-Only Optical Navigation: System Design and Flight Results. *Journal of Guidance, Control, and Dynamics*, 36(6):1576–1595, 2013.

[9] Justin Kruger, Aviad Golan, and Simone D'Amico. Resilient Positioning and Resident Space Object Identification Using Miniaturized Star Tracker Hardware. In *Small Satellite Conference*, Salt Lake City, Utah, 2026.

[10] David Vallado and Paul Crawford. SGP4 Orbit Determination. In *Proceedings of the AIAA/AAS Astrodynamics Specialist Conference*, Honolulu, HI, 2008.

[11] 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.

[12] Justin Kruger and Simone D'Amico. Observability Analysis and Optimization for Angles-Only Navigation of Distributed Space Systems. *Advances in Space Research*, 73(11):5464–5483, 2023.

[13] S. D'Amico. *Autonomous Formation Flying in Low Earth Orbit*. PhD thesis, Delft University, 2010.

[14] K. T. Alfriend, editor. *Spacecraft Formation Flying: Dynamics, Control, and Navigation*. Butterworth-Heinemann/Elsevier, 2010.

[15] 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.

[16] 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.

[17] D. Brouwer. Solution of the problem of artificial satellite theory without drag. *Astronomical Journal*, 64(1274):378–397, 1959.

[18] John A Kennewell and Ba-Ngu Vo. An overview of space situational awareness. In *Proceedings of the 16th International Conference on Information Fusion*, pages 1029–1036, 2013.

[19] Makena Young. Trends in Global Space Situational Awareness. In *Proceedings of the 22nd Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS)*, 2021.

[20] Mariel Borowitz. Examining the Growth of the Global Space Situational Awareness Sector: A Network Analysis Approach. *Space Policy*, 59:101444, 2022.

[21] Mariel Borowitz, Cristina Pérez Hernández, Matt Hejduk, Aurelie Gillet, Monique Moury, and Pascal Faucher. A Technical Comparison of the Public SSA Services in the United States and the European Union. In *Proceedings of the 25th Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS)*, 2024.

[22] Daniel Oltrogge, Joseph Chan, David Vallado, Jeff Cornelius, Andrew D'Uva, and Robert Hall. Deep operator and SSA collaboration for space sustainability. *Journal of Space Safety Engineering*, 11(2):342–361, 2024.

[23] Office of Space Commerce. Traffic coordination system for space (TraCSS), 2026.

[24] U.S. Space Command. Space-track.org: Space surveillance data, 2026.

[25] Tim Flohrer, Holger Krag, Heiner Klinkrad, Benjamin Bastida Virgili, and Carolin Früh. Improving ESA's collision risk estimates by an assessment of the TLE orbit errors of the US SSN catalogue. In *Proceedings of the 5th European Conference on Space Debris*, ESA SP-672, Darmstadt, Germany, April 2009.

[26] David Vallado and Paul Cefola. Two-line element sets - practice and use. *Proceedings of the International Astronautical Congress, IAC*, 7:5812–5825, 01 2012.

[27] Daniel L. O'ltrogge and James A. Cooper. Space Situational Awareness & Space Traffic Management. In Matteo Madi and Olga Sokolova, editors, *Space Debris Peril: Pathways to Opportunities*, chapter 2, pages 9–66. CRC Press, Boca Raton, 2021.

[28] Daniel Oltrogge, Jeff Cornelius, Tim Glinski, David Vallado, Alexander Ridgeway, Yosuke Yamamoto, Salvatore Alfano, Joseph Chan, and Charles Law. Addressing the debilitating effects of maneuvers on SSA accuracy and timeliness. In *Proceedings of the 24th Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS)*, Wailea, Maui, Hawaii, September 2023.

[29] Amelia Bloom, Joshua Wysack, Jacob D. Griesbach, and Anna Lawitzke. Space and ground-based SDA sensor performance comparisons. In *Proceedings of the 23rd Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS)*, Wailea, Maui, Hawaii, September 2022.

[30] Hu Yunpeng, Li Kebo, Liang Yan'gang, and Chen Lei. Review on strategies of space-based optical space situational awareness. *Journal of Systems Engineering and Electronics*, 32(5):1152–1166, 2021.

[31] Justin Kruger and Simone D'Amico. Autonomous navigation of a satellite swarm using inter-satellite bearing angles. *Journal of Guidance, Control, and Dynamics*, 49(4):957–974, 2026.

[32] Laila Kazemi, Alexander Vandenberghe, Mikel Samson, Bram Vandoren, and Tjorven Delabie. Debris detection using star tracker concept verification. In *Proceedings of the 38th Annual AIAA/USU Conference on Small Satellites*, Logan, UT, August 2024.

[33] Randa Qashoa, Matthew Driedger, Ryan Clark, Paul Harrison, Michael Berezin, Regina Lee, and Andrew Howarth. SPACEDUST-optical: Wide-FOV space situational awareness from orbit. In *Proceedings of the 24th Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS)*, Wailea, Maui, Hawaii, September 2023.

[34] Roger C. Hunter, Elwood F. Agasid, Samson Phan, Nathan A. Benz, Scott T. Miller, Caleb A. Adams, Brian C. Kempa, Jeremy D. Frank, Michael J. Iatauro, Justin Kruger, Simone D'Amico, Julianna L. Fishman, and Sasha V. Weston. NASA's starling mission and space situational awareness implications. In *Proceedings of the 26th Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS)*, Wailea, Maui, Hawaii, September 2025.

[35] Scott Miller, Caleb Adams, Nahum Alem, Howard Cannon, et al. Starling CubeSat Swarm Technology Demonstration Flight Results. In *Proceedings of the 38th Small Satellite Conference*, Logan, UT, 2024.

[36] 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.

[37] 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.

[38] Jacco Geul, Erwin Mooij, and Ron Noomen. Tle uncertainty estimation using robust weighted differencing. *Advances in Space Research*, 59(10):2522–2535, 2017.

[39] Justin Kruger and Simone D'Amico. On-Orbit Performance and Lessons Learned for Autonomous Angles-Only Navigation of a Satellite Swarm. In *IEEE Aerospace Conference*, Big Sky, Montana, 2025.

[40] E.A. Wan and R. Van Der Merwe. The unscented kalman filter for nonlinear estimation. In *Proceedings of the IEEE 2000 Adaptive Systems for Signal Processing, Communications, and Control Symposium (Cat. No.00EX373)*, pages 153–158, 2000.

[41] Matthew Willis and Simone D'Amico. Fast Angles-Only Relative Navigation Using Polynomial Dynamics. *Advances in Space Research*, 73(11):5484–5500, 2024.

[42] J. Kruger. *Flight algorithms for autonomous tracking and navigation of distributed space systems using inter-satellite bearing angles*. PhD thesis, Stanford University, 2024.

[43] European Space Agency. DISCOSweb: Database and information system characterising objects in space, 2026.