NASA's Starling mission and space situational awareness implications

Roger C. Hunter, Elwood F. Agasid, Samson Phan, Nathan A. Benz, Scott T. Miller, Caleb A. Adams, Brian C. Kempa, Jeremy D. Frank, and Michael J. Iatauro,

NASA Ames Research Center

Justin Kruger and Simone D'Amico

EraDrive Inc.

Julianna L. Fishman and Sasha V. Weston

Millenium Engineering & Integration Company

ABSTRACT

The Small Spacecraft and Distributed Systems (SSDS) program within NASA's Space Technology Mission Directorate, expands the ability to execute unique missions through rapid development and demonstration of capabilities for small spacecraft applicable to exploration, science and the commercial space sector. Through targeted development and frequent space testing, the program: enables execution of missions at lower cost than previously possible; reduces the time requirement for development of spacecraft; enables new mission architectures through the use of small spacecraft; expands the reach of small spacecraft to new destinations and challenging new environments; and enables the augmentation of existing assets and future missions supporting small spacecraft.

NASA's Starling mission is advancing the readiness of various technologies for cooperative groups of spacecraft – also known as distributed missions, clusters, or swarms. Starling is demonstrating technologies to enable multipoint science data collection by several small spacecraft flying in swarms. The mission uses four CubeSats in low-Earth orbit to test four technologies that let spacecraft operate in a synchronized manner without resources from the ground. Starling launched July 17, 2023 from Rocket Lab Launch Complex 1 in New Zealand. This paper discusses the four technologies carried on Starling with an emphasis on two: Distributed Spacecraft Autonomy (DSA) and Starling Formation-Flying Optical Experiment (StarFOX).

1. SMALL SPACECRAFT & DISTRIBUTED SYSTEMS

Within NASA's Space Technology Mission Directorate (STMD), the Small Spacecraft & Distributed Systems (SSDS) program advances small spacecraft platform capabilities to enable the affordable, rapid development and demonstration of science and exploration missions. Formerly the Small Spacecraft Technologies (SST) program, SSDS leverages the size, weight, power, and cost (SWaP-C) of small spacecraft to advance novel technologies that enable distributed mission architectures. Leveraging the 2024 Civil Space Shortfall Ranking, which identifies current technological gaps in space exploration, SSDS targets key technology investment areas to support future NASA missions, national priorities, and U.S. commercial interests. SSDS is aligned to advance cislunar and Mars activities. Through partnerships, SSDS collaborates with the U.S. space industry, NASA centers, other government agencies, and academia to harness rapid commercial space innovation and apply it to challenging mission needs. Program-funded projects may be carried out at academic institutions, within the commercial sector, or at NASA centers, as public-private partnerships or cooperative agreements.

Small spacecraft are traditionally defined as those with a mass of 180 kilograms or less and capable of being launched into space as an auxiliary or secondary payload. Small spacecraft are not limited to Earth-orbiting satellites but can also include interplanetary spacecraft and later may also serve as planetary re-entry vehicles and landing craft. CubeSats are a special category of small spacecraft and are of particular interest because launch opportunities tend to be more frequent and affordable compared to other small spacecraft, due to the standard sizes and containerization of CubeSats.

Distributed systems are defined as those whose capabilities are physically dispersed across multiple locations, requiring at least two interconnected nodes. These systems execute their mission through either coordinated decision

making—enabled by inter-satellite communication and require interaction between nodes—or distributed decision making, where multiple agents contribute to system level decision. Both approaches greatly enhance space situational awareness (SSA) and space traffic management (STM) activities. Human-centric processes for assessing and acting upon conjunction risks do not scale with the number of satellites being put into orbit. As space becomes increasingly crowded and dynamic, and as interest in understanding the location of orbital objects expand beyond Earth, space-based methods will improve management, awareness and reduce the risk of conjunctions.

2. STARLING MISSION OVERVIEW

One of NASA's proven demonstration missions, Starling set out to advance distributed space technology in the areas of onboard autonomous data collection, relative navigation among spacecraft using optical sensors, and autonomous maneuver planning and execution for STM. Original Starling mission objectives were to demonstrate four novel technologies:

1. Mobile Ad Hoc Networking (MANET)
2. Reconfiguration and Orbit Maintenance Experiments Onboard (ROMEO)
3. Distributed Spacecraft Autonomy (DSA)
4. Starling Formation-Flying Optical Experiment (StarFOX)

Together, these technologies represent key advancements necessary to enable a swarm to accomplish future exploration objectives that require a large sensor network or multi-point science data collection. The successful completion of Starling, now referred to as Starling 1.0, quickly led to follow-on activities known as the Starling 1.5 extension mission. Starling 1.5 prioritized autonomous space traffic coordination, specifically conjunction assessment (CA) and collision avoidance (COLA) of small spacecraft including objectives to demonstrate:

1. Onboard CA for Starling's planned maneuvers;
2. Continuous CA checking of passive and active/maneuvering objects;
3. A ground-based SSA / STM hub that facilitates on-orbit autonomous CA / COLA; and
4. COLA maneuver of Starling spacecraft in response to an onboard CA detection.

In the second mission extension, Starling 1.5+ will validate enhanced capabilities of DSA and StarFOX enabled by new onboard navigation software. Additional scope included:

1. Demonstration of "send and forget" autonomy for swarms;
2. Crosslink ranging for increased intra-swarm situational awareness;
3. Improved space domain awareness capabilities; and
4. Demonstration of surface asset tracking for wildlife migration applications.



Fig. 1. Artist rendition of the Starling swarm in orbit.

The four Starling spacecraft were launched from New Zealand aboard a Rocket Lab Electron launch vehicle, "Baby Come Back", as a secondary payload on July 17, 2023 (see artistic image in Fig. 1). The Starling spacecraft were the first to deploy into a near-sun synchronous orbit at ~575 km altitude; the orbital inclination was determined by the primary payload to be deployed into sun synchronous orbit at a higher altitude. To distinguish between Starling spacecraft, each spacecraft is identified as Starling Vehicle (SV) in the following text.

The following sections describe the successful demonstrations of four Starling technologies: Section 3 combines MANET and ROMEO; and Sections 4 and 5 respectfully capture DSA and StarFOX accomplishments and their follow-on experiments in the Starling 1.5 mission. The DSA and StarFOX experiments are emphasized due to their significant impact on future SSA and STM activities.

3. STARLING COMMUNICATIONS AND CONTROL ACCOMPLISHMENTS

Starling's accomplishments are represented both in the overall success of the mission and in the individual achievements of each novel technology onboard. This section discusses Starling's autonomous networking among the spacecraft and the autonomous planning and execution of maneuvers – such as station-keeping or formation changes. Both MANET and ROMEO offer key capabilities for spacecraft swarms or formation flying missions and further information on both experiments can be found in [1].

3.1 Mobile Ad-hoc Network (MANET)

The MANET experiment on the Starling 1.0 mission demonstrated the first autonomous spacecraft networking in orbit, marking a shift from traditional fixed communication protocols to self-organizing networks. The experiment proved that multiple spacecraft could automatically discover each other and form a communication network without pre-programmed knowledge of their neighbors. This represented a breakthrough for swarm missions where spacecraft configurations change during operations or where communication needs cannot be predetermined before launch.

The system used the BATMAN (Better Approach To Mobile Ad-hoc Networking) protocol integrated with CesiumAstro CommPack radios – software-defined packet radios with dual transmit-receive modules providing near-spherical communication coverage from each spacecraft. The network proved remarkably stable, with spacecraft maintaining consistent direct contact exceeding expectations and making testing multi-hop routing capabilities challenging. Engineers deliberately disabled radios on intermediate spacecraft to force multi-hop scenarios, successfully demonstrating two-hop and three-hop routing which is the first of such demonstrations in space.

Performance of the MANET exceeded all targets. Network establishment, projected to take up to five minutes, was accelerated to under a minute. Network reliability consistently exceeded 90%, and the system validated four different data transmission rates that greatly surpassed original goals. File transfers worked successfully both directly between spacecraft and through intermediate routing nodes. This first-of-its-kind space-based mobile ad hoc networking experiment proved that spacecraft swarms can operate with self-organizing communication networks similar to terrestrial mobile systems, establishing the foundation for more sophisticated and adaptable future space missions.

3.2 Reconfiguration and Orbit Maintenance Experiments Onboard (ROMEO)

The ROMEO experiment set out to demonstrate the ability for the swarm spacecraft to autonomously plan and execute both station-keeping and formation change maneuvers. Autonomous reconfiguration is important for swarms of spacecraft to operate independently, especially for large-scale swarms and those operating in deep space, when communication with ground operators is unreliable.

Much of the demonstration period was spent attempting to tune the onboard navigation filter to get reliable and accurate results, without which, the onboard maneuver planning component could not produce maneuvers that stayed within delta-V limits or had the intended station-keeping effect on the orbit. Eventually, the navigation filter was bypassed altogether, and the maneuver planner used the GPS receiver's reported position and velocity directly for maneuver planning. The maneuver planner itself used an optimization algorithm to minimize delta-V with burns to be completed within a maneuver window specified by ground operators.

After many on-orbit attempts at producing acceptable maneuver plans in November 2024, SV4 executed an autonomous station-keeping maneuver consisting of three burns. This success of accomplishing an autonomous station-keeping maneuver for a single spacecraft still falls short of reliable autonomous maneuvering of a spacecraft swarm but is nevertheless a significant step towards this goal. Future efforts under the Starling 1.5+ mission extension are shifting to a different approach using onboard navigation software provided by NASA's Goddard Space Flight

Center (GSFC). Early results show accurate navigation filter performance, with onboard maneuver planning and execution to commence in the second half of 2025.

4. DISTRIBUTED SPACECRAFT AUTONOMY (DSA)

A primary goal of DSA is to enable autonomous decision-making and execution in multi-spacecraft operations where human operators cannot control the space system in the loop. Previous DSA publications described work in these areas in detail [3][4][5][6] though the technology had not been demonstrated in space until the Starling 1.0 mission. The DSA experiments on Starling 1.0 marked a significant milestone as the first demonstration of fully autonomous distributed operations on multiple spacecraft. The DSA experiment demonstrated enhanced autonomy of distributed spacecraft systems and in practical applications that can be used for science and SSA. To run software consistently across different computing environments, DSA leveraged containerization on the ground for modeling, simulation, validation, and verification.

4.1 The DSA Flight Demonstration

The primary DSA flight demonstration included a GPS Channel Selection Experiment that used a dual-band GPS receiver to autonomously optimize GPS channel selection across the spacecraft swarm for analyzing plasma density in the ionosphere. By analyzing the measurements total electron content (TEC) of the plasma between the spacecraft and GPS satellites, the experiment captured various phenomena in the ionosphere, such as the Equatorial Ionization Anomaly and the Polar Patches [7][8]. In Starling 1.0, TEC events are transient and represent an ideal test case distributed autonomous reaction and control with a spacecraft swarm to perceive and characterize a brief on-orbit event. This GPS Channel Selection Experiment was selected as the primary demonstration due to its ability to showcase autonomous reconfiguration in response to natural phenomena without significant integration efforts or modifications to the spacecraft hardware. Fig. 2 illustrates the GPS Channel Selection DSA Experiment.

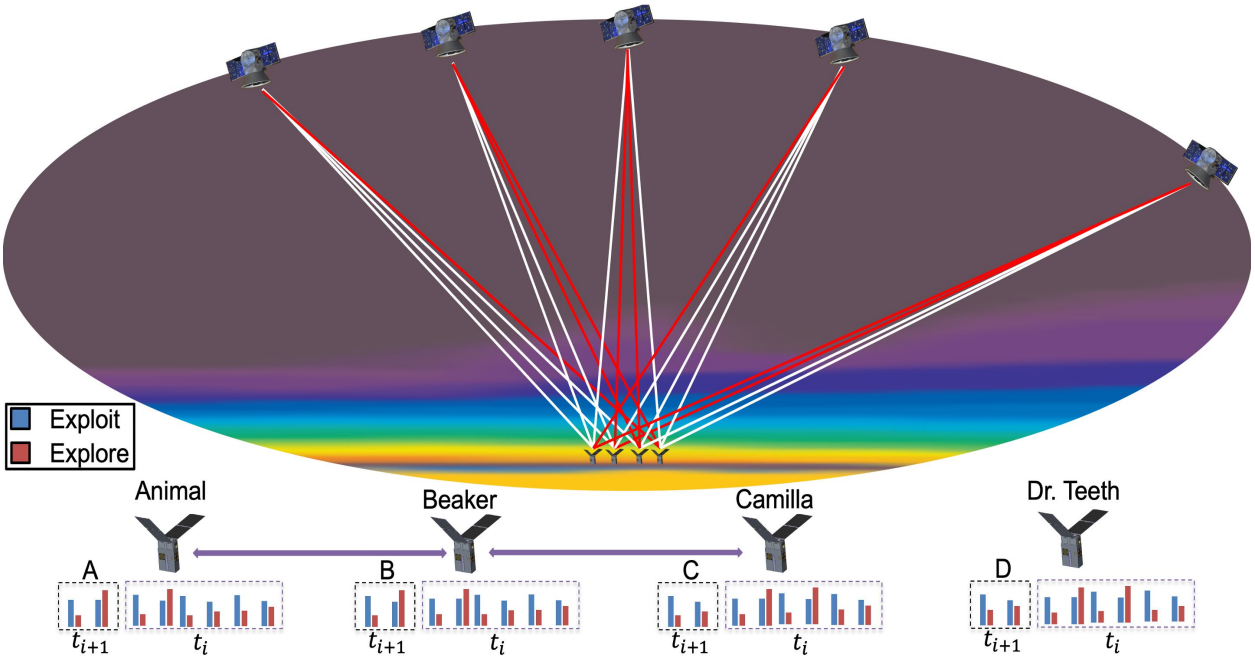


Fig. 2. A graphical representation of the DSA software performing autonomous GPS channel selection observing a transient phenomenon in orbit.

The topside ionosphere is a transitional region between the ionosphere and the inner magnetosphere that displays many dynamic features. By measuring the relative group delay between signals broadcasted at different frequencies by GPS satellites, the dual-band GPS receiver can capture a wide range of ionospheric phenomena. Two specific phenomena of interest, the Equatorial Plasma Bubbles [7] and Polar Patches [9], exhibit distinct behavior in TEC, and thus act as the features to be observed during the experiment.

DSA 1.0 demonstrated the calculation of approximated electron density measurement within the upper ionosphere, specifically relative TEC by listening for signals of opportunity from transmissions through the upper ionosphere, typically these are global navigation satellite system (GNSS) signals. Spikes in electron density are rapid and hard to predict; they make a perfect test case for distributed autonomous reaction and control with a spacecraft swarm and can find ready corollary in SSA. When the observed phenomena were large and homogeneous, the experiment used explorative channel selections— to observe as many channels as possible with a preference for those linked to more distant GPS satellites. In contrast, when the phenomena were spatially constrained and short-lived, the exploitative channel selections focused observations on channels where the TEC count was highest.

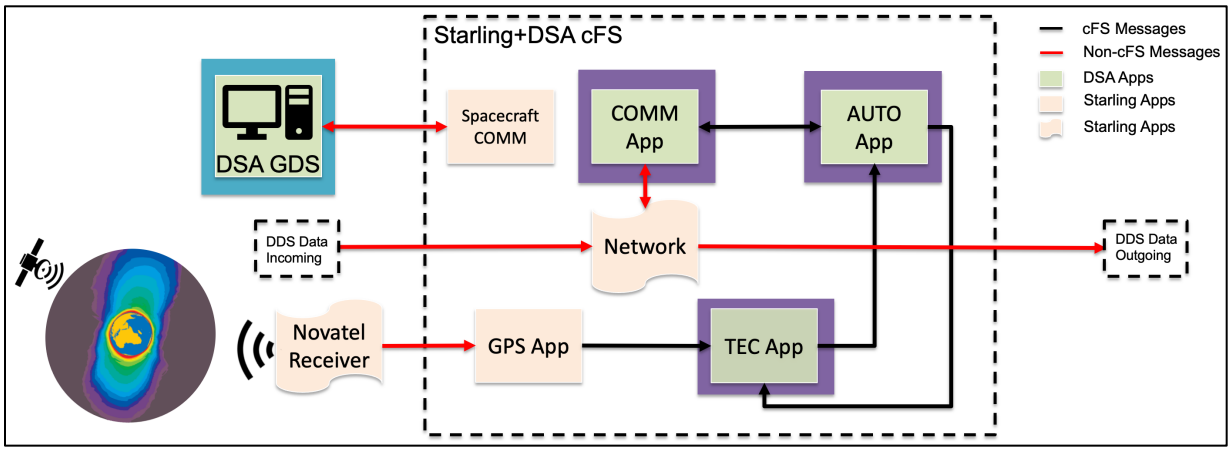


Fig. 3. A simplified diagram of the DSA applications operating within the Starling flight software environment, receiving GPS data, communicating with the DSA ground data system (GDS), and utilizing the DDS network.

Fig. 3. is a simplified representation of a channel assignment scenario within the Data Distribution Service (DDS) network, where multiple spacecraft received signals from GPS satellites. The experiment constrained the number of channels each spacecraft observed, which required DSA to coordinate channel assignments across the network using shared sampling. In the case of spatially-constrained phenomena, simultaneous sampling allowed multiple spacecraft to observe the features of interest from different vantage points. The performance of DSA was evaluated based on the ability to match the optimal channel allocations and the responsiveness to changes in observed features and operating conditions.

4.2 DSA Design

The DSA flight software integrated with the Starling flight software and was built within NASA's core Flight System (cFS) framework. Additionally, Real Time Innovation's Connext DDS Micro framework was used as the communication middleware due to its ability to handle packet-based communication, its scalability, and wide development support. The DSA flight mission software consisted of three applications within the core Flight System framework that included communications (COMM), TEC, and autonomy (AUTO). The COMM app acts as a wrapper for the DDS Micro, enabling message routing over the Ad-Hoc Network of Starling 1.0; the TEC app processes GPS receiver data and provides inputs to the AUTO app, which generates a plan for monitoring GPS channels based on inputs from the local TEC app and other satellites.

The TEC app produced two reward values, Explore and Exploit, that is sent to the AUTO application. The Exploit reward is an approximation of the true TEC that is computed from data from the GPS receivers. The Explore reward is the Euclidean distance between each Starling spacecraft and each GPS spacecraft that an individual Starling spacecraft can view. The AUTO app used a Mixed Integer Linear Programming (MILP) solver to optimize channel allocations by combining rewards from the local TEC app and other Starling spacecraft via the COMM app. The AUTO app then constructed a MILP solver to simultaneously assign GPS channels to all the Starling spacecraft. Each spacecraft was only allowed to monitor a small number of channels, so part of the MILP represents the 'capacity' (number of channels each spacecraft can monitor). Each spacecraft also determined which GPS satellites were in and out of view; these constraints are reflected in the MILP. The MILP also included the objective, be it to Explore or Exploit, depending on configuration. The MILP also contained various weights, allowing blending of Explore and Exploit objectives, and additional constraints such as minimum channel coverage, also subject to configuration. The constraints and objective are also robust to communications dropouts and lack of response from the local TEC app.

Each DSA instance was designed to pose and solve the same channel selection problem using the AUTO app, with the same TEC rewards, that is sent and received across the Starling swarm by the COMM app and BATMAN. Once each AUTO app found the channel assignments to the Starling spacecraft, each DSA instance extracted the 'local' plan from the solution and monitors those channels.

The integrated DSA design offers scalability, fault response capabilities, and seamless integration with the flight software. Notably, all DSA instances compute TEC rewards from GPS, share TEC, assemble and solve the Channel Selection problem, and then monitor the chosen GPS channels, at a rate of 1 Hz.

4.3 DSA On-Orbit: Experiment Periods

For initial DSA experiments, operational products and planning were organized into 24- or 48-hour segments called experiment periods (EP), and DSA had a total of 28 EP opportunities. Early mission operations were challenging, and extensive work was required to build and maintain a stable 4-node platform for the Starling mission. Thus, early attempts to run the DSA experiment failed often. DSA ran on-orbit during 16/28 (57.14%) of EP opportunities and the success rate with orbital operations improved over time. When taking a rolling sum of the last 5 EP deliveries, DSA's success rate reached 80%. A successful DSA run was defined as:

- The DSA mission produced an outcome or delivered a product for execution on orbit.
- The DSA mission successfully ran on orbit and results were delivered.
- This increased the mission's likelihood of meeting full success.

A semi-successful DSA run was defined as:

- The DSA mission missed an opportunity but was not set back according to our operational schedule.
- This did not change the mission's likelihood of meeting full success.

Of the 16 EPs when DSA ran on orbit, 13 of those are considered successful or semi-successful runs. Thus, the total success rate after platform stability was approximately 86.7%. Fig. 4 below shows the EP success rate over time, demonstrating the improved operational capabilities of DSA.

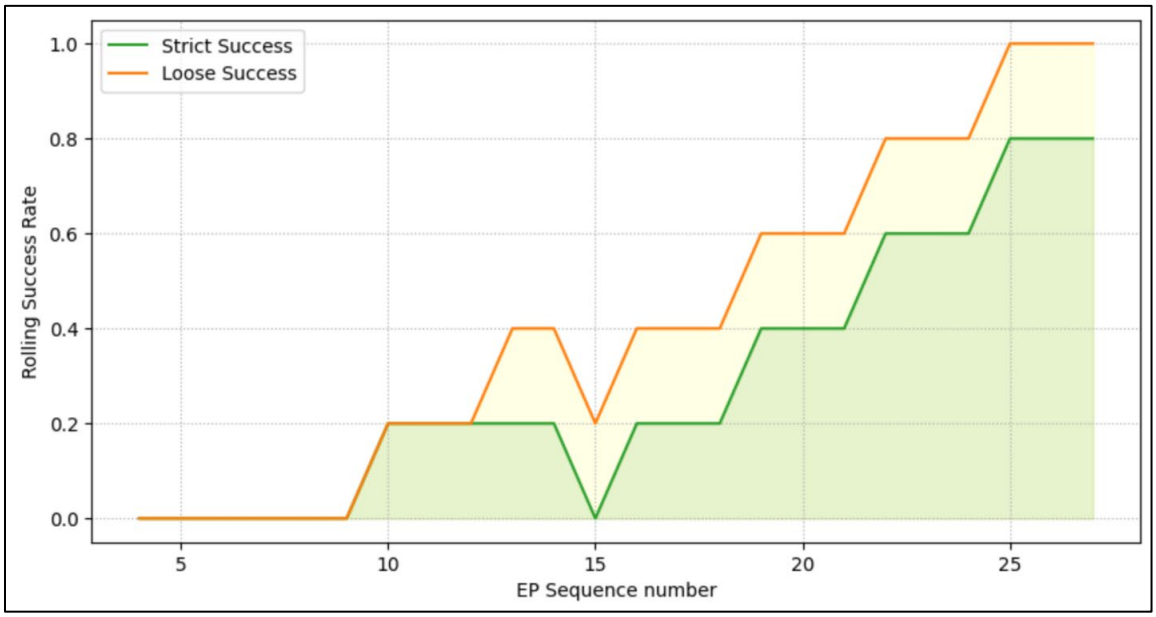


Fig. 4. DSA's EP success rate over time, presented as a rolling average of the prior 5 EPs.

4.4 DSA On-Orbit: Consensus, Coverage and Latency

Recall that each DSA instance was designed to pose and solve the same channel selection problem using the AUTO app, with the same TEC rewards, that is sent and received across the Starling swarm by the COMM app and BATMAN. Once each AUTO app identifies the assignment of channels to Starling spacecraft, each DSA instance extracts the 'local' plan from the solution and monitors those channels. *Consensus* exists when all AUTO apps produce the same plan. Periods of non-consensus occur due to changes in network topology and configuration, GPS satellite

visibility set changes, changes in Explore or Exploit reward between communication and MILP setup, or changes in the AUTO app configuration. Fig. 5 shows consensus during an experiment period.

SV	Time	SV2 plan	SV3 plan	SV4 plan
SV2	01:19:03	(16, 17, 27, 30)	(8, 9, 14, 21)	(3, 4, 6, 7)
SV3	01:19:03	(16, 17, 27, 30)	(8, 9, 14, 21)	(3, 4, 6, 7)
SV4	01:19:03	(16, 17, 27, 30)	(8, 9, 14, 21)	(3, 4, 6, 7)
...	...	...	...	...
SV2	01:19:06	(16, 17, 27, 30)	(8, 9, 14, 21)	(3, 4, 6, 7)
SV3	01:19:06	(16, 17, 27, 30)	(8, 9, 14, 21)	(3, 4, 6, 7)
SV4	01:19:06	(16, 17, 27, 30)	(8, 9, 14, 21)	(3, 4, 6, 7)
SV2	01:19:07	(16, 17, 27, 30)	(8, 9, 14, 21)	(3, 4, 6, 7)
SV3	01:19:07	(16, 17, 27, 30)	(9, 14, 21, **22**)	(4, 6, 7, 8)
SV4	01:19:07	(16, 17, 27, 30)	(8, 9, 14, 21)	(3, 4, 6, 7)

Fig. 5. A sequence of 4 consecutive periods of consensus at 2024-03-22 01:19:03 - 01:19:06, followed by no consensus at 01:19:07 (indicated by bold font).

Coverage requires all GPS satellites in view were selected by at least one plan generated by the AUTO app. Coverage is desirable when Explore is the objective. Coverage is not possible if too many GPS satellites are in view, or alternatively, the capacity of Starling satellites to monitor GPS channels is too small. Fig. 6 shows coverage during an experiment period.

SV	SV2 plan	SV3 plan	SV4 plan
SV2	(17, 21, 27, 30)	(7, 8, 9, 14)	(2, 3, 4, 6)
SV3	(17, 21, 27, 30)	(7, 8, 9, 14)	(2, 3, 4, 6)
SV4	(17, 21, 27, 30)	(7, 8, 9, 14)	(2, 3, 4, 6)

SV	Visible GPS Satellites
SV2	(2, 3, 4, 6, 7, 8, 9, 14, 17, 21, 27, 30)
SV3	(2, 3, 4, 7, 8, 9, 14, 17, 21, 27)
SV4	(2, 3, 4, 6, 7, 8, 9, 14, 17, 21, 27, 30)
All	(2, 3, 4, 6, 7, 8, 9, 14, 17, 21, 27, 30)

Fig. 6. Full coverage and consensus at 2024-03-22 01:16:08; GPS visibility sets displayed in All row.

One definition of *latency* is the difference between the time an event occurs that requires reconfiguration, and the time that consensus is achieved. The same events that lead to a lack of consensus may lead to the need to reconfigure, and thus allow measurement of latency. However, it is possible that numerous events occur in rapid succession, and reconfiguration may not always be needed even when events do occur. Fig. 7 shows an example of reactive operations that illustrates latency. In this case, multiple GPS satellites leave the visibility sets of DSA spacecraft at different times. However, after a 2-tick period of stability of GPS visibility sets, DSA was able to reach consensus. Fig. 8 shows that while consensus may be lost, the time needed to re-establish consensus is often quite small (seconds).

SV	Time	SV2 plan	SV3 plan	SV4 plan
SV2	01:14:58	(7, 21, 27, 30)	(7, 8, 9, 14)	(3, 4, 6, 22)
SV3	01:14:58	(17, 21, 27, 30)	(7, 8, 9, 14)	(3, 4, 6, 22)
SV4	01:14:58	(7, 21, 27, 30)	(7, 8, 9, 14)	(3, 4, 6, 22)
SV2	01:14:59	(7, 21, 27, 30)	(7, 8, 9, 14)	(3, 4, 6, 22)
SV3	01:14:59	(**17**, 21, 27, 30)	(7, **9**, 14, **22**)	(3, 4, 6, 8)
SV4	01:14:59	(7, 21, 27, 30)	(7, 8, 9, 14)	(3, 4, 6, 22)
SV2	01:15:00	(7, 21, 27, 30)	(8, 9, 14, 22)	(3, 4, 6, 7)
SV3	01:15:00	(7, 21, 27, 30)	(8, 9, 14, **17**)	(3, 4, 6, **22**)
SV4	01:15:00	(**17**, 21, 27, 30)	(8, 9, 14, 22)	(3, 4, 6, 7)
SV2	01:15:01	(7, 21, 27, 30)	(8, 9, 14, **17**)	(3, 4, 6, 22)
SV3	01:15:01	(**17**, 21, 27, 30)	(**7**, 8, 9, 14)	(3, 4, 6, 22)
SV4	01:15:01	(7, 21, 27, 30)	(8, 9, 14, **17**)	(3, 4, 6, 22)
SV2	01:15:02	(7, 21, 27, 30)	(7, 8, 9, 14)	(3, 4, 6, 22)
SV3	01:15:02	(17, 21, 27, 30)	(7, 8, 9, 14)	(3, 4, 6, 22)
SV4	01:15:02	(7, 21, 27, 30)	(7, 8, 9, 14)	(3, 4, 6, 22)

SV	Time	Visible GPS Satellites
SV2	01:14:58	(2, 3, 4, 6, 7, 8, 9, 14, 17, 21, 27, 30)
SV3	01:14:58	(2, 3, 4, 6, 7, 8, 9, 14, 17, 21, 27)
SV4	01:14:58	(2, 3, 4, 6, 7, 8, 9, 14, 17, 21, 22, 27, 30)
SV2	01:14:59	(2, 3, 4, 6, 7, 8, 9, 14, 17, 21, 27, 30)
SV3	01:14:59	(2, 4, 8, 9, 14, 17, 21)(**3**, **6**, **7**, **27**)
SV4	01:14:59	(2, 3, 4, 6, 7, 8, 9, 14, 17, 21, 22, 27, 30)
SV2	01:15:00	(2, 3, 4, 6, 7, 8, 9, 14, 17, 21, 27, 30)
SV3	01:15:00	(2, **3**, 4, **7**, 8, 9, 14, 17, 21)
SV4	01:15:00	(2, 3, 4, 6, 7, 8, 9, 14, 17, 21, 22, 27, 30)
SV2	01:15:01	(2, 3, 4, 6, 7, 8, 9, 14, 17, 21, 27, 30)
SV3	01:15:01	(2, 3, 4, 7, 8, 9, 14, 17, 21, **27**)
SV4	01:15:01	(2, 3, 4, 6, 7, 8, 9, 14, 17, 21, 22, 27, 30)
SV2	01:15:02	(2, 3, 4, 6, 7, 8, 9, 14, 17, 21, 27, 30)
SV3	01:15:02	(2, 3, 4, 7, 8, 9, 14, 17, 21, 27)
SV4	01:15:02	(2, 3, 4, 6, 7, 8, 9, 14, 17, 21, 22, 27, 30)

Fig. 7. Reactive operations during the period 2024-03-22 01:14:58 - 01:15:02. Consensus is lost at 2024-03-22 01:14:59, then achieved at 2024-03-22 01:15:02, as shown in the left table. GPS visibility sets are shown in the

right; bold font shows newly visible GPS, the second parenthetical set indicates GPS satellites are no longer visible at 01:14:59. GPS satellite visibility is consistent from 01:15:01 - 01:15:02.

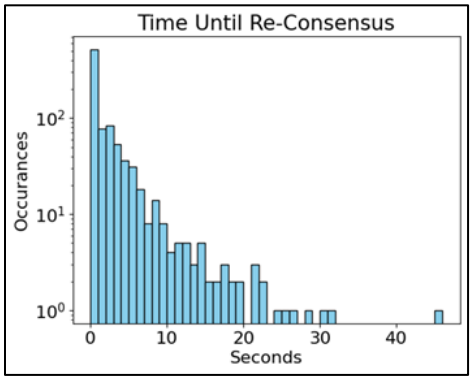


Fig. 8. Time until re-consensus over the selected period of 01:14:57 - 02:50:38 in experiment period Sequence Number 25, re-consensus happens quickly and often (log scale) with a worst-case outlier of 46 seconds.

4.5 DSA On-Orbit: Key Performance and Mission Firsts

At the completion of Starling 1.0, DSA accomplished several firsts-in-space [10] including:

- First fully distributed autonomous operation of multiple spacecraft;
- First use of space-to-space communications to autonomously share state information between multiple spacecraft;
- First demonstration of fully distributed reactive operations onboard multiple spacecraft;
- First use of fully distributed automated planning onboard multiple spacecraft; and
- First use of a general-purpose automated reasoning system onboard a spacecraft.

During the original mission, DSA and Starling 1.0 consistently established a 3-spacecraft cross-link network via the DSA COMM application, laying the foundation for future DSA communication systems that could build upon the BATMAN protocol and DDS network layer [1][11]. The successful demonstration of this network topology was crucial for the DSA experiment, as it enables the autonomous decision-making and distributed file sharing capabilities essential for multi-spacecraft operations. Finally, Fig. 9 shows DSA's Key Performance Parameters (KPPs), indicating DSA successfully accomplished the threshold objectives in reducing data uplink, and more than met its time to reconfiguration objectives. The complete methodology for measuring these KPPs is described in [11].

DSA Key Performance Parameters				
KPP	SOA	Threshold	Goal	In-Flight Result
data uplink reduction	50%	66%	75%	49.3 % - 61.3 %
Time to reconfiguration	NA	22 min	5 min	2.04 - 6 sec

Fig. 9. Distributed Spacecraft Autonomy's Key Performance Parameters

4.6 DSA 1.5: Follow-on Experiments and Results

In subsequent Starling 1.5 experiments, DSA 1.5 focused on improving its science proxy towards a higher-fidelity data product computed at the edge, as well as addressing emergent autonomy needs from Starling 1.0, particularly local autonomy and robust crosslink data transfer.

The purpose of improving the TEC measurement was to demonstrate that high-quality observation data could be generated onboard comparable to data typically produced through ground-based data processing systems, see Fig. 10. Performing this on-orbit allows the autonomy to react to better quality observations. For TEC, this is done by smoothing the pseudorange values calculated by the GPS receiver and correcting for differential code bias that gets introduced by both the transmitters on board the GPS satellites and the receiver on the Starling spacecraft. DSA 1.5 was able to incorporate those corrections as well as a measure of data quality, both of which are available for the autonomy to incorporate in its decision-making.

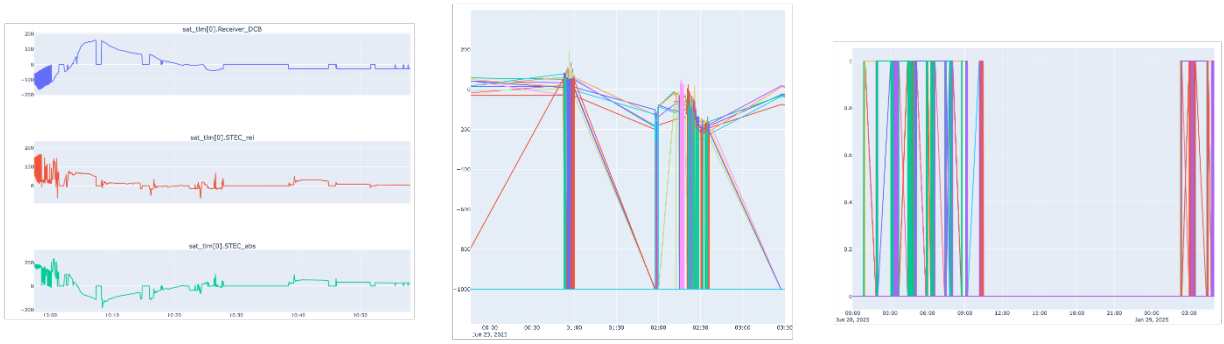


Fig. 10: From right to left, (1) a decomposition of receiver code bias through absolute total electron content (TEC) calculation, (2) absolute TEC measurements, (3) TEC measurement quality.

The autonomy focus of DSA 1.5 was to develop advanced capabilities for distributed space systems by increasing the robustness, resilience and self-sufficiency of spacecraft. While DSA successfully demonstrated autonomous distributed decision-making onboard Starling 1.0, this was accomplished only in software. To demonstrate the ability to command spacecraft hardware subsystems (e.g., crosslink radios) rather than just re-configuring software would greatly increase confidence in onboard autonomy. Similarly, the ability to safely and autonomously create, and power on and off, the cross-link network during operations will raise mission efficiency by decreasing power consumption, reducing bandwidth, and overall reducing spacecraft mission operation planning.

DSA was able to autonomously create, break down, and re-create the crosslink network in response to changes in the radio frequency (RF) environment—TEC was again used as a proxy for demonstration purposes—as well as orbit position. Additional DSA successes included onboard decision-making enabled by NASA Ames' Plan Execution Interchange Language (PLEXIL) Executive agent. This agent controlled spacecraft operations in scenarios requiring human in-the-loop intervention. While PLEXIL has a long history of supporting remote robotics for Earth-based analog missions simulating off-Earth environments (e.g., the Moon, Mars, and Europa), this marks its first deployment in space.

Finally, DSA 1.5 demonstrated the ability to send large files over the Ad-Hoc crosslink network by redelivering failed file segments. This redelivery was based on standard transfer measures in the COMM application and an additional layer of autonomy. DSA 1.5 file delivery (known as LFT) was event-driven instead of pre-scheduled based on absolute or relative time, thus more robust operations were enabled.

The motivation for an autonomous, distributed LFT system comes from Starling 1.0 operations where one of the spacecraft proved to have a severe anomalous limit on the ground uplink rate that was restricted to approximately 10 KB per day. Since the Starling DSA flight software binaries were between 2MB - 8MB (depending on the deployed applications), this would have required approximately 27 months to upload. Transfers using the crosslink radios were orders of magnitude faster and not limited by the dynamics of a spacecraft moving relative to a fixed ground station.

DSA's COMM app is approximately 3MB, and thus we choose to define a large file as one that is between 3MB and 10MB. For the mission extension we defined the minimum success for a LFT transfer as 3MB and full success as 10MB (max order of magnitude). The LFT app is inspired by peer-to-peer file sharing, with the objective of facilitating reliable and secure distributed file sharing between any node on the network. What distinguishes LFT from traditional centralized file-sharing systems is its decentralized architecture, which allows for multiple nodes to send file fragments. Furthermore, LFT introduces a feature that enables nodes to 'forcibly' upload or push files to other nodes. LFT was successfully tested on the DSA 1.5 mission extension, during which the spacecraft successfully and repeatedly transferred 10MB files with guarantees on integrity and validity across multiple network interruptions lasting 60 minutes.

5. STARLING FORMATION-FLYING OPTICAL EXPERIMENT (STARFOX)

The primary mission phase of the StarFOX experiment began in December 2023 and ended in May 2024 and represents the first demonstrations of autonomous angles-only navigation for a satellite swarm for multiple target objects and multiple observers [13]. In addition to this technology ‘first’, three additional specific advances over the

state of the art for autonomous navigation for distributed space systems (DSS) were achieved on orbit during this mission. These include:

1. First autonomous initialization of relative navigation for unknown target objects;
2. First long-term convergence of angles-only orbit estimates without requiring maneuvers; and
3. First instance of simultaneous absolute and relative orbit determination using angles-only measurements, without GNSS access.

The goal of StarFOX was to eliminate GNSS dependencies and exploit the Starling swarm to enhance the robustness and autonomy of optical navigation, while building upon the successes and lessons learned of prior missions. Compared to prior angles-only flight experiments, such as the Advanced Rendezvous using GPS and Optical Navigation (ARGON) [14] and the Autonomous Vision Approach Navigation and Target Identification (AVANTI) [15], StarFOX presents a significantly more autonomous, flexible, and robust demonstration of navigation in flight. Operational experience has also provided valuable lessons learned for the design and implementation of vision-based navigation aboard future swarms and DSS. The success of StarFOX led to the 'StarFOX+' follow-on experiment Starling 1.5, which extends the StarFOX flight software with new capabilities for on-orbit SSA activities. StarFOX+ operations commenced in March 2025. Flight results from StarFOX and StarFOX+ are described in the following section.

5.1 Angles-Only Navigation

During angles-only navigation, observer spacecraft use onboard vision-based sensors (VBS) to obtain the inter-satellite bearing angles to target space objects. Bearing angle measurements are typically expressed as azimuth α and elevation ε in the camera frame and are used to estimate the absolute state of the observer as well as the relative states of targets with respect to the observer. If multiple cooperative observers (i.e., satellites) are present with an inter-satellite link (ISL), the observers can share their individual observations directly without relying on ground stations. By sharing these bearing angle measurements with each other, observer spacecraft can combine their data to create a more accurate and comprehensive view of the object's position and movement, known as distributed stereovision. Fig. 11 provides a schematic example of swarm geometry, images processed on-board, and interpretation of bearing angles.

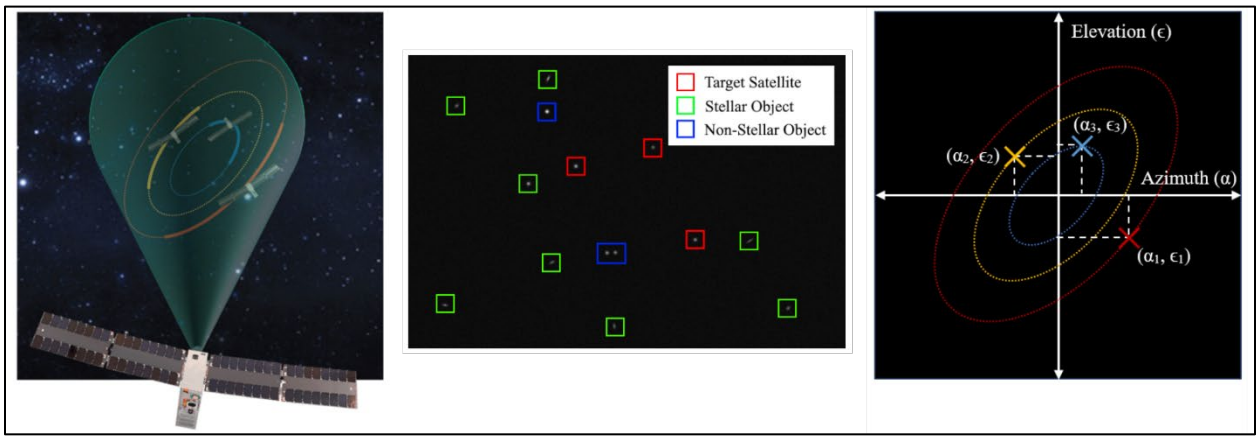


Fig. 11. Notional schematic of angles-only navigation for one observer and three targets: satellite geometry (left), the camera image (center), and derived azimuth and elevation (right).

Angles-only technologies are therefore compelling in the context of autonomy, logistics and SSA. Onboard VBS reduce reliance on external metrologies such as GNSS or ground station tracking. Cameras themselves are low size-weight-power-cost sensors ubiquitous on modern spacecraft in the form of star trackers.

5.2 ARTMS Navigation Architecture

StarFOX applies the angles-only Absolute and Relative Trajectory Measurement System (ARTMS), a self-contained architecture that provides distributed, autonomous, scalable navigation for DSS orbiting an arbitrary central body, without reliance on maneuvers or external measurement sources [18]. A high-level overview of ARTMS as implemented in the StarFOX flight software is shown in Fig. 12. Each active observer in the DSS runs an identical ARTMS instance to navigate for itself, other observers, and any detected targets.

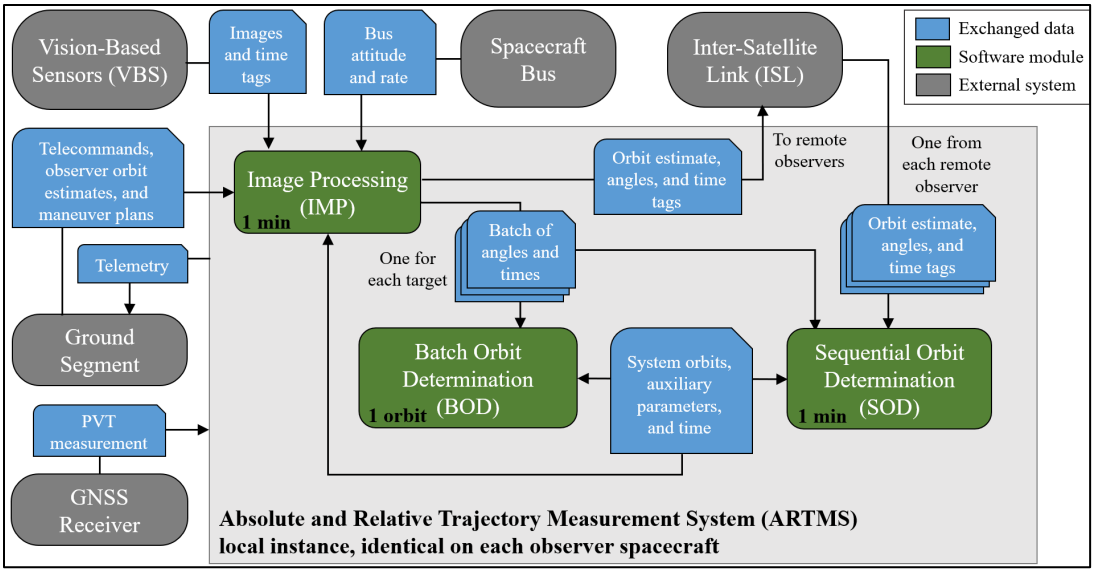


Fig. 12. A high-level overview of the Absolute and Relative Trajectory Measurement System flight architecture.

ARTMS consists of three core modules: IMage Processing (IMP), Batch Orbit Determination (BOD), and Sequential Orbit Determination (SOD). The VBS provides ARTMS with time-tagged images which are processed to obtain inertial bearing angles to resident space objects (RSO). The ISL communicates orbit estimates and bearing angles between observers, which allows ARTMS to perform distributed navigation. The ground segment provides telecommands, maneuver plans and orbit estimates to each observer, and receives ARTMS telemetry. Position/Velocity/Time (PVT) navigation solutions may optionally be provided by a GNSS receiver.

The IMP module uses VBS images from the onboard camera to produce batches of bearing angles and corresponding uncertainties to all visible targets [19]. IMP measurement batches are provided to the BOD and SOD modules, and to the ISL for transmission. The BOD module uses angle batches from IMP to compute orbit estimates for itself and all observed targets [20]. These orbit-state estimates are provided to the SOD for initialization, fault detection, and initialization of the SOD navigation filter that continuously refines the observer's absolute orbit estimate, target relative orbit estimates, and auxiliary state estimates (e.g., empirical accelerations) [21]. This filter seamlessly fuses local measurements from the IMP and remote measurements from the ISL. SOD estimates are provided to the IMP to assist with onboard target tracking.

Overall, ARTMS provided 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 (preferably a star tracker to refer measurements to the inertial frame, but not necessarily) and an optional ISL for cooperative navigation. No external measurement sources were needed, and each observer only required a coarse estimate of its own absolute orbit at a single epoch to initialize. Future ARTMS versions will support entirely lost-in-space optical positioning with on-orbit testing planned for Q4 2025.

5.3 StarFOX Flight Results: Autonomous Swarming

In its original form, StarFOX proposed four primary navigation modes, as enabled by ARTMS:

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

Modes 1 and 3 are most applicable to cooperative swarm scenarios in Earth orbit; Mode 2 is most useful for SSA in which unknown targets must be detected and tracked; and Mode 4 is most useful for deep space applications in which GNSS is unavailable. In total, 18 experiment blocks were executed over 34 experiment days. Eleven StarFOX experiment goals were defined prior to launch, summarized in Table 1. All thresholds and the majority of goals were accomplished through flight experiments or post-processing of flight data [13]. The achieved relative position accuracies of tens of meters in the R/N directions and hundreds of meters in the T direction are sufficient to support a variety of swarm objectives, and the level of autonomy demonstrated in flight is a strong enabler for future distributed missions.

Table 1. Summary of StarFOX performance goals and flight results.

#	Objective	Threshold	Goal	Result
1	Tracking multiple targets simultaneously	1 target	3 targets	2 targets (in flight)
2	Accuracy of relative position knowledge with one observer	Estimate produced	1% error relative to ISD	0.5% error relative to ISD (in flight)
3	Convergence time for relative position knowledge with one observer	Convergence observed	6 orbits	2 orbits (in flight)
4	Accuracy of relative position knowledge with multiple observers	Estimate produced	0.1% error relative to ISD	0.5% (in flight) 0.1% (on ground)
5	Convergence time for relative position knowledge with multiple observers	Convergence observed	2 orbits	1 orbit (in flight)
6	Accuracy of relative position knowledge in presence of maneuvers	Estimate produced	0.1% error relative to ISD	Estimate produced (in flight)
7	Convergence time for relative position knowledge in presence of maneuvers	Convergence observed	2 orbits	Convergence observed (in flight)
8	Accuracy of absolute position knowledge with multiple observers	Estimate produced	1 km error	850 m error (on ground)
9	Convergence time for absolute position knowledge with multiple observers	Convergence observed	24 hours	Convergence observed (on ground)
10	Time to produce autonomous relative orbit initialization	Initialization produced	2 orbits	5 orbits (in flight)
11	Accuracy of target range estimate for relative orbit initialization	Initialization produced	20% error relative to ISD	16% (in flight) 2% (on ground)

5.4 StarFOX+ Flight Results: On-Orbit Space Situational Awareness

For the follow-on Starling 1.5 mission, StarFOX+ endowed ARTMS with five new navigation capabilities [22]. The two capabilities most relevant for SSA are described below.

1. *Unknown RSO identification in IMP.* Identify unknown RSO detected in images by matching them to known objects from a space object catalog. This capability enables the collection of bearing angle measurements of RSOs by space-based optical sensors. Measurements may then be used to update the observer's onboard RSO catalog or downlinked to improve the accuracy and timeliness of global SSA products.
2. *Fast cooperative orbit initialization in BOD*. Apply multi-observer methods to reduce the number of measurements required to generate an initial relative orbit estimate in BOD. This enables fast initialization of orbit estimates for newly-tracked RSOs, on a timescale of minutes. New RSOs may then be added to existing catalogs to improve their completeness

5.4.1 RSO Identification

Presented here are the initial flight results from April – July 2025, which are the first on-orbit demonstrations of autonomous identification and tracking of RSOs by an observer satellite performed completely on-board.

It is assumed the observer satellite has access to an onboard catalog of RSOs, which are given as two-line elements (TLE) at the catalog epoch. When a new image arrives, catalog orbits are analytically propagated to the image epoch using the SGP4 model. Propagated TLEs are converted to a position-velocity state and rotated into the camera reference frame, such that their predicted bearing angles can be computed. If the bearing angle lies within the field of view and the inter-satellite distance (ISD) is less than a user-specified maximum, the corresponding bearing angle covariance is computed using the empirical model defined in [22].

Before each flight experiment, the most recent RSO catalog was downloaded from Space-Track [23], and relevant data for target identification was extracted. The experiment team prepared telecommands that were integrated with the catalog and subsequently uplinked to the spacecraft. An identification was considered correct if $e_{geo} / d_{RSO} \leq 0.001$,

or incorrect if e_{geo} / $d_{RSO} \leq 0.01$, where d_{RSO} is the ISD; all other identifications were considered ambiguous. The measurement noise floor of the star tracker was approximately 20-40" (1σ) [24], which corresponds to e_{geo} / d_{RSO} = 0.0001.

Examples of flight images reconstructed from detected regions of interest are provided in Fig. 13 and illustrate on-orbit detections of 1) a Starling swarm satellite; 2) the defunct Centre National d'Etudes Spatiales (CNES) SPOT-1 satellite; and 3) a Starlink satellite. The farthest object identified was North American Aerospace Defense Command (NORAD) ID 40112, a Long March CZ-4C rocket booster with an ISD of 2997 km, on June 27, 2025. The closest object observed was Starling 4 with an ISD of 64 km, on June 19, 2025. Across the experiment set, 591 unique objects were identified. The most frequently observed object, aside from Starling swarm members, was the PACE satellite (NORAD ID 58928), which was identified in 26 images.

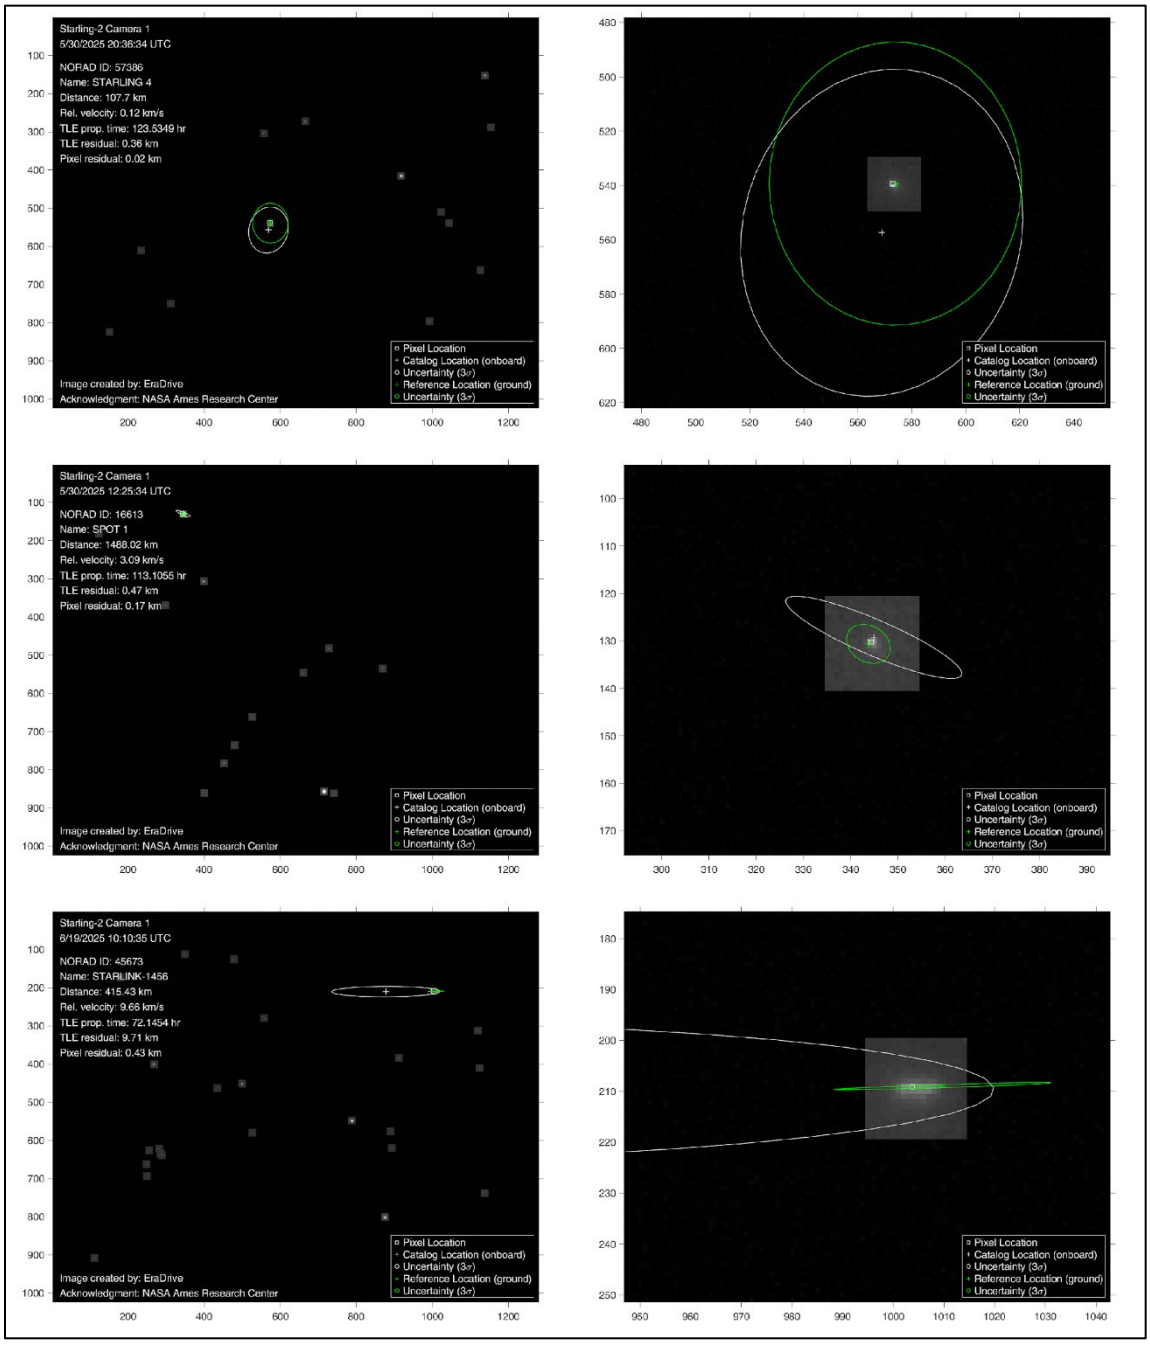


Fig. 13. Reconstructed flight images (left) with the identified object enlarged (right).

Table 2 presents statistics regarding on-orbit RSO identifications. Across all results, the correct match rate is 81%, which is 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. The Starling category possesses the lowest ISD, but their covariance ellipses are largely parallel with the camera boresight, meaning the size of the ellipse remains bounded. The Starlink category often features large components of relative velocity perpendicular to the camera boresight which produces 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 σ_{assign} and d_{max} thresholds.

Table 2. RSO identification and measurement statistics.

	Correct Match	Ambiguous Match	Incorrect Match	Correct Match Rate	e_{ang} (Mean $\pm 1\sigma$, correct matches only)	e_{geo}	e_{cat}
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

Table 2 also presents mean angular, geometric and catalog residuals. Angular residuals are generally 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 significant differences in fidelity between Starlink TLE solutions, as propagated onboard, and Starlink predicted ephemeris data, as used for reference truth on the ground. Fig. 14 presents histograms of geometric residuals across all identifications, where Starling residuals are generally lowest and Starlink residuals are highest.

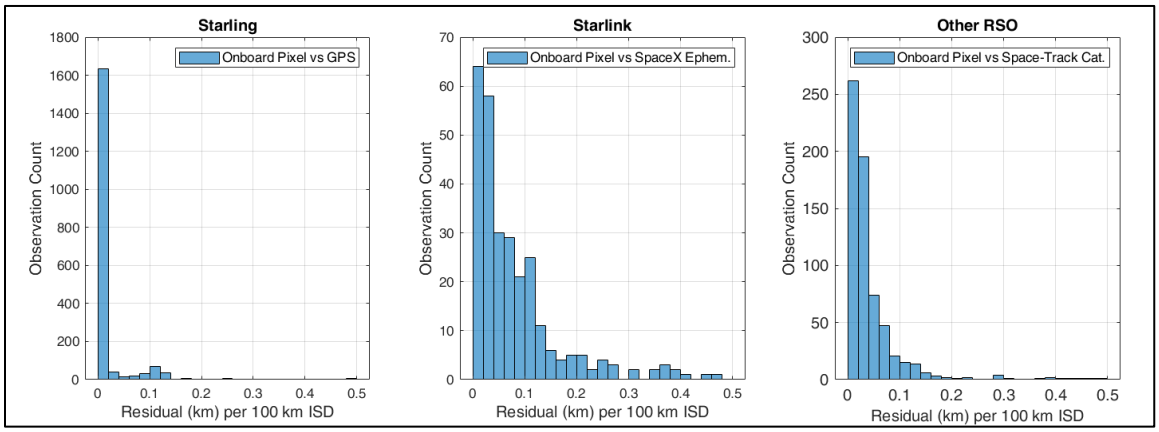


Fig. 14. Geometric residuals between the measured bearing angle and reference position.

The truest picture of measurement accuracy is provided by Starling measurements since their reference truth is derived from GPS (likely the most accurate reference among those surveyed). In Table 2, the mean e_{geo} from on-orbit measurements is approximately 14 m per 100 km of ISD which approaches the level of GPS positioning accuracy for nearby targets. Furthermore, most non-Starling measurements in Fig. 14 achieve errors less than 100 m per 100 km of ISD. Flight results indicate that on-orbit optical sensing is a valuable measurement source for improving SSA accuracy, given that the positioning accuracy of current TLE products is typically much larger (ranging from hundreds of meters to several kilometers) [25][26].

5.4.2 Fast Orbit Initialization

Assume a batch of line-of-sight (LOS) vectors from observer satellite to target *i* is obtained by the onboard camera of observer *j*. Following the reasoning of Willis [27], each LOS vector between *i* and *j* is parallel to the relative position vector of *i* with respect to *j*. Taking the cross product therefore leads to a system of constraint equations. The relative state of target *i* is characterized here by the 6D vector of quasi-nonsingular relative orbit elements (ROE) [28], and there exists a linear mapping between the ROE and target relative position in RTN. This mapping can be substituted into the constraint equations to form a linear system, to be solved for the unknown ROE. For a single observer and

linear mapping, the ROE solution can be scaled by any amount and still be valid. However, when using two observers with a known baseline, the location of the common target is fixed by the intersection of their LOS vectors. The new cooperative BOD system combines the LOS batch from the onboard camera with the LOS batch from second observer *k*, to form an extended linear system to solve for the ROE of target *i*. To produce an uncertainty estimate and simultaneously perform averaging for the effects of measurement noise on the solution, a sampling approach is used. Consider randomly selecting N_b measurements from a larger batch where $N_b \geq 3$. The resulting linear system can be solved to produce a solution sample. The process can be repeated to generate N_s state samples, such that the mean and standard deviation of the sample set provides the output ROE mean and uncertainty.

On-orbit tests for the fast initialization procedure are scheduled for August 2025. To emulate on-orbit behavior, flight imagery and GPS data from March 24, 2024 are passed through the ARTMS C++ flight software running in a MATLAB Simulink environment. In the experiment scenario, SV3 and SV4 track SV2 while in a passive safety ellipse (PSE) formation. The VBS and GPS sample times are 120 seconds, and the ISL messages are broadcasted every 120 seconds between SV3 and SV4. The maximum allowed timespan of the combined measurement batch is limited to 10 minutes, i.e., up to 6 bearing angles per observer, with $N_b = 3$ and $N_s = 100$. In comparison, the single-observer BOD system for StarFOX requires 50-100 measurements across 1-2 orbits.

Fig. 15 visualizes one relative orbit initialization in the δa-$\delta\lambda$, δe_x-δe_y, δi_x-δi_y ROE planes. Errors in δi are less than 50 m; errors in δe are less than 200 m; δa are less than 500 m; and errors in $\delta\lambda$ are less than 2 km. Errors remain within 3σ covariance bounds and $\delta\lambda$ errors fall well below the 20% of ISD threshold defined for StarFOX pre-flight [18]. Estimation performance is sufficient to initialize target tracking and sampling adequately characterizes uncertainty.

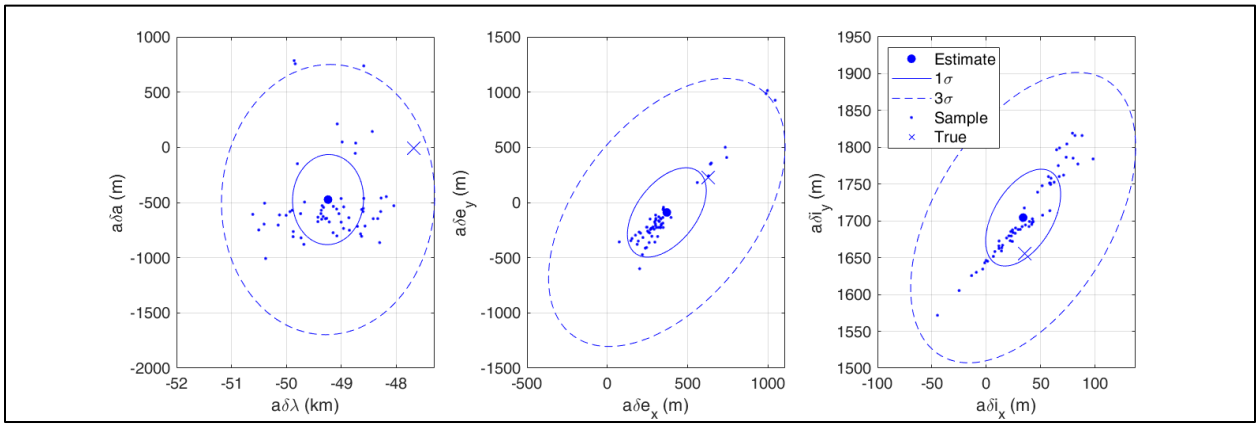


Fig. 15. Relative orbit elements error, uncertainty, and samples for one initialization of Starling Vehicle (SV) 2 from SV4.

Table 3 presents statistics for the first eight cooperative initializations during the simulation. Sufficient target visibility is obtained after 4200 seconds, after which an initialization is computed every 10 minutes. Performance varies significantly between trials. The maximum $\delta\lambda$ error is 18 km (approximately 35% of target range) which is poorly captured by the computed state uncertainty from sampling. The minimum $\delta\lambda$ error is 2.5 km (approximately 5% of target range), which is captured within the 3σ uncertainty bounds.

Table 3. BOD estimation errors and 1σ uncertainties for each initialization.

Time (s)	$a\delta a$ (m)	$a\delta\lambda$ (km)	$a\delta e_x$ (m)	$a\delta e_y$ (m)	$a\delta i_x$ (m)	$a\delta i_y$ (m)	Resid. (rad)
4200	-634 ± 169	-10.29 ± 0.27	-109 ± 123	-307 ± 120	128 ± 125	370 ± 118	0.0217
4800	-381 ± 818	-8.39 ± 1.38	-36 ± 349	-85 ± 752	-42 ± 58	112 ± 151	0.0222
5400	-297 ± 572	-2.52 ± 1.08	-236 ± 246	-131 ± 526	-36 ± 11	35 ± 36	0.0209
6000	-366 ± 511	-11.55 ± 0.97	15 ± 216	-163 ± 472	-30 ± 76	162 ± 194	0.0232
6600	-487 ± 416	-5.18 ± 0.79	28 ± 175	-67 ± 383	-5 ± 42 m	67 ± 111	0.0215
7200	-471 ± 416	-18.19 ± 1.96	-87 ± 397	1 ± 884	-44 ± 129 m	255 ± 340	0.0283
8400	-302 ± 477	-2.61 ± 0.84	-136 ± 204	-193 ± 439	-22 ± 26 m	18 ± 66	0.0243
9000	-421 ± 427	-5.89 ± 0.84	80 ± 184	-37 ± 392	-31 ± 44 m	49 ± 116	0.0245

The tested scenario is particularly challenging given the short measurement period, large time interval between images, and significant measurement noise onboard. The swarm scenario also gives rise to poor angles-only observability due to the similarity of observer and target orbits. It is therefore promising that even under these conditions, reasonable state initializations can be computed. Results are especially impactful for SSA scenarios in which targets may have high relative velocities and spend little time in the visual range of the sensor. Additionally, for many SSA applications, targets are in dissimilar orbits to the observer. The resulting diversity of measurement geometry and rapid target motion greatly improves angles-only observability when compared to swarm case, with corresponding improvements to estimation performance expected.

6. STUDIES

In addition to on-orbit demonstrations of critical technologies for future STM and SSA architectures, SSDS sponsored studies to guide future investments. Two current studies are described below.

6.1 On-orbit Space Management System

Towards advancing the role of small spacecraft and distributed systems for space applications, SSDS partnered with NASA Ames Research Center to conduct a study examining the technical approaches, feasibility, trade space, and limitations of using small spacecraft for STM efforts. This builds upon the success of the overall Starling mission to evaluate the utility of optical observations for SSA, aligning with numerous Civil Space Shortfall Rankings (1589, 1262, 1432, 147, and 1438 respectively). The focus of this study is to (1) reduce the uncertainty in orbit determination for the various classes of observable objects of interest in space (debris, other spacecraft); (2) track, characterize, enable timeliness of data delivery for actionable decision making; and (3) identify ground systems and/or on-orbit resources needed for these determinations. The results of the study will be distributed to parties of interest within NASA, the Department of Commerce, and the Department of Defense, for informing future space-based monitoring architectures.

6.2 Next Generation Swarm

The Next-Generation Multi-Agent Swarm (NGS) study is being conducted by NASA Ames to develop a comprehensive understanding of emerging multi-agent swarm capabilities in cislunar and lunar space. The study aims to identify existing swarm capabilities and assess their potential for persistent lunar SSA, surface monitoring, and distributed autonomy demonstrations. A key objective is to design a next-generation swarm that can perform autonomous distributed remote sensing, position, navigation, and timing (PNT), and automated deployment that leverages autonomy, edge computing, and interoperable networking to enable cooperative operations without the need for immediate human assistance. This study will address specific shortfalls, including 1625, 1438, 1433, 1557, 1431, and 1430.

7. CONCLUSIONS AND FUTURE WORK

This paper discussed the notable accomplishments of NASA's Starling mission which benefit several areas of space exploration. It is remarkable how many firsts this mission has achieved in only two years in space— and continues to operate at the time of writing. The SSDS Starling 1.0 and extended 1.5 flight demonstration results highlight successes in the areas of onboard autonomous data collection, relative navigation among small spacecraft using optical sensors, and autonomous maneuver planning and execution for space traffic management.

At the completion of Starling 1.0 in 2024, all four novel technologies had achieved their objectives. MANET demonstrated autonomous communication using an Ad-Hoc network between multiple spacecraft; ROMEO illustrated autonomous spacecraft cluster maintenance within a swarm; StarFOX accomplished autonomous, GPS-independent optical navigation; and DSA achieved autonomous distributed operation of multiple spacecraft. The results of the StarFOX and DSA experiments quickly led to the StarFOX+ and DSA 1.5 demonstrations on the Starling 1.5 follow-on mission that focused on SSA and STM activities.

The demonstration of GPS-independent tracking methods by StarFOX are especially notable for missions in environments where GPS is unavailable, such as cislunar and lunar space, deep space, or in certain condensed Earth orbits. Optical relative navigation has been received with high interest in the small spacecraft community primarily as exploration goes where GPS or other traditional navigation aids are lacking. StarFOX+ demonstrated the first on-board detection and tracking of RSOs using a commercial-off-the-shelf (COTS) star-tracker. To date, 591 unique

RSOs up to an inter-satellite distance of 2997 km have been detected in flight thus far, with optical measurement residuals approaching GPS accuracy. This capability greatly benefits SSA and STM as there are no current cost-effective solutions to these expanding issues. The demonstrations by StarFOX and StarFOX+ to track and detect space objects with high angular resolution is momentous for SSA and STM efforts as it improves awareness of relative positions and motions of space objects.

DSA's demonstrated ability to reconfigure swarm formations or trajectories without external commands reduces reliance on ground-based control, a crucial step for managing large satellite constellations. Its GPS Channel Selection experiment allowed Starling 1.0 to demonstrate several autonomous swarm operations that are critical to supporting an on-orbit SSA capability, including autonomous state sharing, distributed reasoning, automated planning, and reactivity to space events. Reconfiguration can also improve the spacecraft's responsiveness in dynamic orbital environments, such as areas with heavy space traffic, where it can detect and respond to potential real-time anomalies. Additionally, by sharing the orbital state information via space-to-space communication, the community is able to have improved collective awareness. This adaptation to SSA and STM can finally enable technological solutions to the overcrowded space problem. In the 1.5 extension mission, DSA 1.5 demonstrated on-orbit AI software for small spacecraft planning, scheduling, and execution that successfully demonstrated the first use of their local-autonomy agent to control the swarm, enabling the spacecraft to dynamically command themselves. It further demonstrated edge computing of science measurements and adaptive planning and scheduling of cross-link network.

SSDS will continue to invest in demonstrations of new, next generation swarm technologies that enable coordinated, automated SSA and STM operations. Future developments will be aligned with both Civil Space Shortfall Rankings in addition to NASA's directive for an increased application to cislunar space and towards Mars. Future on-orbit demonstrations will feature small satellite swarms for complementary positioning, navigation, and timing and SSA applications. Part of SSDS' charter is to develop innovative ideas that meet the federal government's specific research and development needs with the potential for commercialization. NASA's spacecraft autonomy and SSA technologies developed through the Starling mission has led to successful commercialization follow-on activities. EraDrive, Inc. is a startup spun out of the Space Rendezvous Lab of Stanford University, that focuses on commercializing the StarFOX suite of technologies.

8. ACKNOWLEDGEMENTS

Starling is funded by the SSDS program within NASA's Space Technology Mission Directorate (STMD). Fnding for DSA was provided by Game Changing Development (GCD) program within STMD. The technology underlying StarFOX received development funding under the University SmallSat Technology Partnerships (USTP).

The authors and Starling mission team would like to thank the following for their contributions to the Starling 1.5 mission: SpaceX, the University of Texas at Austin, and NASA's Conjunction Assessment Risk Analysis team. The authors and team would also like to acknowledge the following hardware partners vital to making the mission a success: Blue Canyon Technologies (Lafayette, Colorado), Rocket Lab USA (Long Beach, California), Stanford University Space Rendezvous Lab (Stanford, California), York Space Systems (Denver, Colorado), formerly Emergent Space Technologies, CesiumAstro (Austin, Texas), and L3Harris Technologies, Inc (Melbourne, Florida).

9. REFERENCES

[1] S. Miller, C. Adams, N. Alem, et al., "Starling CubeSat Swarm Technology Demonstration Flight Results," Proceedings of the 38th Small Satellite Conference, Logan, Utah, 2024.

[2] C. Adams, B. Kempa, W. Vaughan, et al. "Development of a High-Performance, Heterogenous, Scalable Test-Bed for Distributed Spacecraft," 2023 IEEE Aerospace Conference, doi: 10.1109/AERO55745.2023.10115695.

[3] D. Cellucci, N. Cramer, and J. Frank, "Distributed Spacecraft Autonomy ," 2020 AIAA Ascend Conference

[4] N. Cramer, D. Cellucci, C. Adams, et al., "Design and Testing of Autonomous Distributed Space Systems," Proceedings of the AIAA/USU Conference on Small Satellites.

[5] C. Adams, B. Kempa, M. Iatauro, et al., "An Overview of Distributed Spacecraft Autonomy at NASA Ames," Proceedings of the AIAA/USU Conference on Small Satellites.

[6] C. Adams, B. Kempa, W., Vaughan, et al., "Development of a High-Performance, Heterogenous, Scalable Test-Bed for Distributed Spacecraft," 2023 IEEE Aerospace Conference, doi: 10.1109/AERO55745.2023.10115695.

[7] N. Balan, L. Liu, and H. Le, "A Brief Review of Equatorial Ionization Anomaly and Ionospheric Irregularities," Journal of Earth and Planetary Physics, Vol. 2, No. 4, 2018, pp. 257-275.
[8] M. Noja, C. Stolle, J. Park, et al., "Long-term analysis of ionospheric polar patches based on CHAMP TEC data," Journal of Radio Science, Vol. 48, No. 3, 2013, pp. 289-301.
[9] A. Chartier, C. Mitchell, and E. Miller, "Annual Occurrence Rates of Ionospheric Polar Cap Patches Observed Using SWARM", Journal of Geophysical Research: Space Physics, Vol. 123, No. 3, 2018, pp. 2327-2335
[10] C. Adams, J. Frank, S. Gridnev, et al., "Advancing Autonomy In Distributed Space Systems: Insights From On-Orbit Testing With The Starling 1.0 Mission," Proceedings of the 2024 Small Satellites Systems and Services Symposium, March 2025.
[11] C. Adams, J. Frank, S. Gridnev, et al., "Advancing Autonomy in Distributed Space Systems: Results from the Distributed Spacecraft Autonomy Experiment on Starling 1.0," Proceedings of the AIAA/USU Conference on Small Satellites, Logan, UT, 2024.
[12] C. Adams, B. Kempa, W. Vaughan, et al., "Development of a High-Performance, Heterogenous, Scalable Test-Bed for Distributed Spacecraft," 2023 IEEE Aerospace Conference, pp. 1-8, 10.1109/AERO55745.2023.10115695.
[13] J. Kruger, S. Hwang, and S. D'Amico, "Starling Formation-Flying Optical Experiment: Initial Operations and Flight Results," Proceedings of the 38th Small Satellite Conference, Logan, Utah, 2024.
[14] S. D'Amico, J. Ardaens, G. Gaias, et al., "Noncooperative Rendezvous Using Angles-Only Optical Navigation: System Design and Flight Results," Journal of Guidance, Control, and Dynamics, Vol. 36, No. 6, 2013, pp. 1576–1595.
[15] G. Gaias and J. Ardaens, "Flight Demonstration of Autonomous Noncooperative Rendezvous in Low Earth Orbit," Journal of Guidance, Control, and Dynamics, Vol. 41, No. 6, 2017, pp. 1337–1354.
[16] D. Woffinden and D. Geller, "Observability Criteria for Angles-Only Navigation," IEEE Transactions on Aerospace and Electronic Systems, Vol. 45, No. 3, 2009, pp. 1194–1208.
[17] J. Kruger and S. D'Amico, "Observability Analysis and Optimization for Angles-Only Navigation of Distributed Space Systems," Advances in Space Research, Vol. 73, No. 11, 2023, pp. 5464–5483.
[18] J. Kruger, A. Koenig, and S. D'Amico, "The Starling Formation-Flying Optical Experiment (StarFOX): System Design, Experiment Design and Pre-Flight Verification," Journal of Spacecraft and Rockets, 2023, pp. 1–23.
[19] J. Kruger and S. D'Amico, "Autonomous angles-only multitarget tracking for spacecraft swarms," Acta Astronautica, Vol. 189, 2021, pp. 514–529.
[20] A. Koenig and S. D'Amico, "Observability-Aware Numerical Algorithm for Angles-Only Initial Relative Orbit Determination," 2020 AAS/AIAA Astrodynamics Specialist Conference, Lake Tahoe, CA, 2020.
[21] J. Sullivan, A. Koenig, J. Kruger, et al., "Generalized Angles-Only Navigation Architecture for Autonomous Distributed Space Systems," Journal of Guidance, Control, and Dynamics, Vol. 44, No. 6, 2021, pp. 1087–1105.
[22] J. Kruger and S. D'Amico, "On-Orbit Performance and Lessons Learned for Autonomous Angles-Only Navigation of a Satellite Swarm," IEEE Aerospace Conference, Big Sky, Montana, 2025.
[23] Space-Track. https://space-track.org/.
[24] S. Palo, G. Stafford, and A. Hoskins., "An Agile Multi-Use Nano Star Camera for Constellation Applications," Proceedings of the Small Satellite Conference, Logan, UT, 2013.
[25] M. Holzinger and M. Jah, "Challenges and Potential in Space Domain Awareness," Journal of Guidance, Control, and Dynamics, Vol. 41, No. 1, 2018, pp. 15–18.
[26] H. Flohrer, H. Krag, H. Klinkrad, B. Virgili, and C. Fruh, "Improving ESA's collision risk estimates by an assessment of the TLE orbit errors of the US SSN catalogue," European Conference on Space Debris, Darmstadt, Germany, 2009.
[27] M. Willis and S. D'Amico, "Fast Angles-Only Relative Navigation Using Polynomial Dynamics", Advances in Space Research, Vol. 73, No. 11, pp. 5484-5500 (2024).
[28] S. D'Amico. "Autonomous Formation Flying in Low Earth Orbit". PhD thesis, Delft University, 2010.