SSC24-I-06

Starling CubeSat Swarm Technology Demonstration Flight Results

Scott Miller, Caleb Adams, Nahum Alem, Howard Cannon, Randy Grashuis, Ted Hendriks, Soon Hwang, Michael Iatauro and Craig Pires
NASA Ames Research Center
M/S 213-13, Moffett Field, CA 94035; 650-495-5320
scott.t.miller@nasa.gov

Justin Kruger and Simone D'Amico (StarFOX PI)
Stanford University
496 Lomita Mall, Stanford, CA 94305; 650-272-9968
jjkruger@stanford.edu

ABSTRACT

The Starling swarm of four 6U CubeSats launched in July 2023 to test four key technologies to enable future swarm missions:

1) Mobile Ad-Hoc Networking (MANET) over a crosslink radio network
2) Autonomous onboard decision-making for operations
3) Optical-based absolute and relative navigation
4) Autonomous maneuver planning and execution

The Starling team implemented the Better Approach to Mobile Ad-hoc Networking (B.A.T.M.A.N.) protocol to automatically manage the crosslink network of four satellites. The B.A.T.M.A.N. protocol uses a decentralized approach to managing a multi-hop mesh network of devices, in this case, a satellite swarm. The four satellites were able to successfully establish a network at multiple data rates and demonstrate file transfer and command issuance between spacecraft over the network.

Starling incorporated Distributed Spacecraft Autonomy's (DSA) software to demonstrate onboard decision-making. The DSA software takes L1/L2 band GPS measurements and uses them to estimate the relative Total Electron Count (TEC) in the ionosphere. The onboard software then determines if there are any features of interest and provides that information to the other satellites over the crosslink network. The swarm of satellites then reaches a consensus on the optimal TEC observation strategy and adjusts its measurement collection tactics autonomously.

The Starling Formation-Flying Optical Experiment (StarFOX), produced by Stanford's Space Rendezvous Laboratory, uses onboard star trackers to collect images of the other swarm spacecraft and produce angles-only orbit estimates for navigation. This system is envisioned to be valuable when Global Navigation Satellite Systems (GNSS) are unavailable, such as for navigation in cis-lunar or deep space or tracking of non-cooperative resident space objects. StarFOX successfully applied its algorithms to navigate multiple spacecraft targets simultaneously, using star tracker imagery.

Finally, Starling used Emergent Space's Cluster Flight Application (CFA) software suite for the Reconfiguration and Orbit Maintenance Experiments Onboard (ROMEO) demonstration of autonomously planning and executing propulsive maneuvers. Large swarms will need to be able to maintain formation requirements with minimal operator involvement, especially as the size of the swarm scales up. Results from the ROMEO experiment are presented.

Starling is funded by the Small Spacecraft Technology (SST) program out of NASA's Space Technology Mission Directorate (STMD).

MISSION OVERVIEW

Purpose

The Starling technology demonstration was conceived to build upon previous Small Spacecraft Technology (SST) missions like Edison Demonstration of Smallsat Networks (EDSN) and Nodes that began to investigate network architectures and dynamic reconfiguration of the network between multiple small spacecraft equipped with crosslink radios.[1,2] In addition to implementing more advanced and scalable networking features, Starling also set out to advance swarm technology in the areas of onboard autonomous science data collection, relative navigation among spacecraft using optical sensors, and autonomous maneuver planning and execution. Together, these four technologies represent key advancements necessary to enable a swarm, or cooperative free-flying system of distributed spacecraft, to accomplish future exploration concepts that require a large sensor network or multi-point science data collection.[3]

Swarm Description

The Starling spacecraft consist of the following major subassemblies:

- XB1 6U CubeSat spacecraft bus, provided by Blue Canyon Technologies (BCT), consisting of two “back-to-back” star trackers, reaction wheels, deployable solar panels, GPS receiver, Globalstar beacon radio, S-Band space-to-ground radio, avionics and power control, propulsion system controller, thermal control system, and mechanical structure. The bus components (aside from the solar panels) fit into a 2U volume, with 4U allocated to the Payload and Propulsion subsystems.
- Payload, consisting of a power and data controller board (custom, by NASA Ames), a Xiphos Q7S space-rated processor, and the CommPack crosslink radio provided by CesiumAstro.
- Propulsion subsystem: the Hollow Additively Manufactured Lightweight Efficient Thruster (HAMLET) cold gas system by NASA Ames.

Figure 1: Model of a single 6U Starling spacecraft

NASA personnel integrated and tested each of the four spacecraft at Ames facilities in preparation for integration into their dispensers.

Figure 2: Four spacecraft integrated into the Rocket Lab Maxwell dispensers

Launch and Early Operations

The spacecraft were launched from New Zealand aboard a Rocket Lab Electron launch vehicle as a secondary payload on the “Baby Come Back” launch on July 17, 2023. The Starling spacecraft were the first to deploy on the mission into a near-sun synchronous orbit (the orbital inclination was determined by the primary payload to be deployed into sun synchronous orbit at a higher altitude) at ~575km altitude. The spacecraft drifted apart for six weeks while spacecraft commissioning activities took place. Upon completion of the commissioning phase, three of the four spacecraft maneuvered to initiate a two-month drift into the “in-train” formation. The swarm operated for four months in the in-train formation. During this phase, the swarm tested the four swarm technology payloads while performing periodic maneuvers to maintain the swarm. After the in-train phase, each spacecraft maneuvered to transition into a passively safe configuration at the same in-track spacing. The swarm technology payload operated for two months in the passively safety ellipse (PSE) formation.

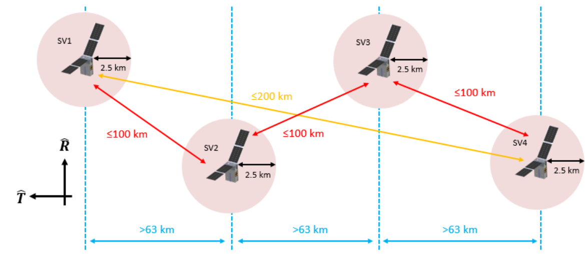

Figure 3: Experiment and spacecraft operational requirements provide primary drivers behind Starling formation design

The combination of the swarm geometry and range requirements for Starling drove the design of the Starling formation. The mission design requirements imposed two distinct formation phases on the swarm: an in-train phase and a phase in passively safe ellipses (PSEs). Relative orbital elements (ROEs), as articulated by D'Amico and Montenbruck[4] were used to design both formations. The orbit control strategy in this paper was also implemented for each formation, due to its ability to maintain Starling's desired separation between spacecraft and to minimize collision hazard. ROEs are defined to describe the motion of the deputy relative to the chief in the Radial-Tangential-Normal (RTN) frame. The first five ROEs define the shape of the cyclical motion in the RTN frame while the sixth ROE defines the in-track spacing along the tangential direction. The two-dimensional vector defined by δe_x and δe_y is referred to as the **eccentricity** vector. The two-dimensional vector defined by δi_x and δi_y is referred to as the **inclination** vector. The passively safe condition (in PSE1 and PSE2) is achieved when these two vectors are either parallel or anti-parallel in the E-I (eccentricity-inclination) plane.

The first maneuver of the mission was the Drift Control Maneuver (DCM), which controlled the post-deployment drift to initiate ingress to the in-train formation. The original DCM approach was to maneuver all four spacecraft to create the in-train formation. Following a propulsion leak on Starling-1, which lowered its orbit with respect to the other spacecraft, a new DCM strategy was needed in order to reduce its >3000 km separation from the others. Due to initial failed attempts of the DCM (due to various spacecraft issues), and Starling-1's inability to maneuver at the time, the modified approach was to maneuver Starlings-2, -3, and -4, to reduce their distance with respect to Starling-1. The phasing period following the DCM was also increased, to slow the drift rates as the spacecraft approached each other. The desired in-train separation was achieved in mid-November.

Table 1: Mission Phases Timeline

Mission Phase	Start Date	Duration
Launch and Orbit	July 17, 2023	45 min
Commissioning	July 17, 2023	6 weeks
Formation Ingress	August 30, 2023	2 months
In-Train Formation	November 14, 2023	4 months
Passive Safety Ellipse	March 7, 2024	2 months

As Starlings-2, -3, and -4 approached Starling-1 following the DCM, they performed weekly station-keeping maneuvers to maintain an inter-satellite separation of 60-100 km. In early March, Starlings-2, -3, and -4 performed the Swarm Reconfiguration Maneuver (SRM) to establish the PSE formation, which introduced radial and cross-track separation between the spacecraft. Although the desired radial and cross-track separation was not fully achieved with the SRMs, no corrections were needed because the offsets achieved met the experiment formation requirements. The spacecraft stayed in this PSE formation for the rest of the mission.

Mission Operations

The NASA Starling swarm mission employed a novel approach to mission operations by dividing responsibilities between Blue Canyon Technologies (BCT) and NASA Ames Research Center. BCT was responsible for satellite bus command and control, bus commissioning, and ground contact scheduling with KSAT-Lite. NASA Ames handled Flight Dynamics, Payload Operations for the four technology experiments, as well as systems engineering and payload software development.

BCT was responsible for delivery of bus telemetry binaries and payload data files to NASA Ames for further processing. The Ames Ground Data System processed the bus telemetry and payload data received from BCT and distributed the processed data to the Experiment teams, Flight Dynamics, and Operations engineers for analysis.

The payload operations process at NASA Ames involved several key steps. First, the Ames team received command input products from the experiment and Flight Dynamics teams. They then generated executable command sequences for on-orbit operations. To ensure the validity and safety of these command sequences, the team conducted command validation using a Hardware-in-the-Loop ground simulator before uploading them to the satellites.

The Starling mission demonstrated the ability to operate with lean staffing levels across various teams. The Flight

Dynamics team at NASA Ames consisted of fewer than 2 full-time equivalents (FTEs), while the Payload Operations team was managed with fewer than 4 FTEs. Additionally, the Systems Engineering and payload software development team operated with fewer than 4 FTEs. Despite these lean teams, they successfully carried out their respective responsibilities, contributing to the overall success of the mission.

The collaboration between BCT and NASA Ames allowed each organization to focus on its area of expertise, resulting in efficient mission operations and successful execution of the technology experiments. This public-private partnership is the first of its kind for a NASA Ames CubeSat mission, showcasing an innovative approach to mission operations that could serve as a model for future missions.

MOBILE AD HOC NETWORKING (MANET)

Technology

Crosslink networking is considered a supporting technology for the "upper" level applications and experiments (StarFOX, ROMEO, and DSA). From a basic level, the Starling objective was to establish an ad hoc network between the spacecraft. Previous missions had used fixed protocols and messages. For Starling, fewer restrictions were imposed on the upper-level applications and they could rely on the ad hoc network protocol to handle traffic routing. This was important as the configuration of the swarm or the scheduled need would not be known prior to deployment and may change during the mission.

B.A.T.M.A.N. protocol – Several MANET protocols were evaluated before the selection. Major ones included OLSR, DSDV, Babel and BATMAN. BATMAN, or Better Approach To Mobile Ad hoc Networking, was chosen due to its maturity, availability and existing implementation into the Linux distribution that Starling used. BATMAN also has the ability to adapt and scale in an environment of limited network traffic and changing topology.[3] No direct modifications of the BATMAN protocol were required, but TCP/IP virtual devices needed to be written to provide an interface between the BATMAN driver software and the crosslink radio.

CesiumAstro CommPack Radio – The CesiumAstro CommPack Radio is a software-defined packet radio. The Crosslink radio is composed of a Software-defined radio (SDR) and two Transmit Receive Modules (TRM) that are both amplifiers and antennas. The TRMs are on either side of the spacecraft to give a reasonable attempt at near-spherical coverage. Within the SDR, processing of signals is performed to packetize data and perform advanced processing such as carrier-sense multiple access (CSMA) to help to deconflict the network traffic.

NASA's core Flight System (cFS) is a foundational piece of the flight software for Starling. cFS applications were developed to communicate and control the CesiumAstro radio for crosslink and provide drivers to interface to the CommPack. Implementation of software consisted of direct interfaces to control the modes of the CesiumAstro radio, interfacing and translating packets into the BATMAN network. Application-level interfaces were also developed to provide commands and telemetry to and from the radio.

Flight Results

The basic objectives of MANET for Starling were to test crosslink communications and the networking software. Tests were conducted to test basic connectivity as well as networking capabilities. Networking included both point-to-point communications as well as routing through spacecraft nodes. Key metrics were crosslink data rate, packet routing speed, time to establish network and network reliability.

Overall, the crosslink network met or exceeded all its key performance parameters. Operational characteristics were discovered that helped optimize the network. For instance, the "normal" time for all nodes to discover each other could be several minutes when only BATMAN messages were being exchanged. The spacecraft always discovered each other within the 5 minute goal, but the time could be accelerated by forcing traffic across the network. This allowed the network to be established in less than a minute.

Another interesting discovery was that the network was almost too reliable to be able to test the network routing with more than two spacecraft, as typically the spacecraft were all within direct radio range of all other spacecraft and could communicate point-to-point. The overall network reliability was typically much greater than 90%. To test routing, experiments were set up to turn off one of the two TRMs on the middle spacecraft in the formation. With this method, both two-hop and three-hop network routing was demonstrated. Overall routing through all spacecraft still was faster than the goal speed and was able to be routed typically under a second.

File transfers were a bit more mixed. Files were successfully able to be sent and received both point-to-point and routed through an intermediate spacecraft by the BATMAN protocol. The radio is capable of running at four speeds between 244 and 1953 ksym/sec. All speeds were tested and proven to work (greatly exceeding the KPP goals). The issue was that files greater than 50 Kbytes proved to be unreliable in being properly received. A few different protocols and methods were attempted, but the 50 KB limit persisted. Future work during an expected mission extension will include software updates for more reliable transfer of larger files over the network.

RECONFIGURATION AND ORBIT MAINTENANCE EXPERIMENTS ONBOARD (ROMEO)

Technology

The Reconfiguration and Orbit Maintenance Experiments Onboard (ROMEO) experiment runs a software payload designed to autonomously maintain the swarm in the desired formation. The base software package for the experiment was provided by York Space Systems (formerly Emergent Space Technologies) and is known as the Cluster Flight Application (CFA) suite, originally developed for DARPA's System F6.[5] It is a modular system that consists of the navigation service, which determines where the satellite is, the orbital maintenance service, which creates maneuver plans to keep the swarm in the desired formation, and the module maneuver service which generates the maneuver commands. To ensure that autonomous maneuvers maintain a safe space environment, autonomously generated maneuvers were to be screened for conjunctions on the ground and a final Go/No Go command was to be given by the ground. Additional constraints were placed on the maneuver planning service, where the total maneuver plan delta-v could not exceed 0.6 m/s for normal station-keeping, and each individual maneuver delta-v could not exceed 0.292 m/s.

Flight Results

Each experiment block was planned to fit within a 48-hour window, where the first half was dedicated to maneuver plan generation and conjunction screening, and the second half dedicated to maneuver plan execution. Maneuver plans would be generated one hour prior to the first ground contact of the day and then downlinked at the next ground pass. These plans are then converted into the required screening products and the Go/No Go decision is made by the ground. If the maneuvers are cleared to proceed, the commands required to allow the software to execute the maneuver are uploaded 24 hours post maneuver plan generation.

ROMEO was unable to plan and execute maneuvers to meet the full success criteria. Issues in the navigation filter prevented the software from obtaining an accurate orbit estimate, and the maneuver command generation process proved unreliable. Experiment runs during December 2023 and January 2024 showed slow convergence and the navigation filter was unable to remain converged for the remainder of the experiment, which interfered with the maneuver planning process. The nav filter was tuned, health checks were adjusted, and the method by which GPS ephemeris messages were delivered improved navigation filter convergence and stability.

A maneuver plan was first generated on February 1st, 2024, but it failed to meet the delta-v restrictions, with a total plan delta-v of 4.56 m/s. Subsequent runs were able to reduce the total delta-v down to approximately 1.20 m/s, but still failed the individual maneuver delta-v threshold. In March, maneuver plans stopped being generated. While investigating the failure to generate maneuver plans, it was discovered that the navigation filter had an error when computing the inclination and right ascension of the ascending node.

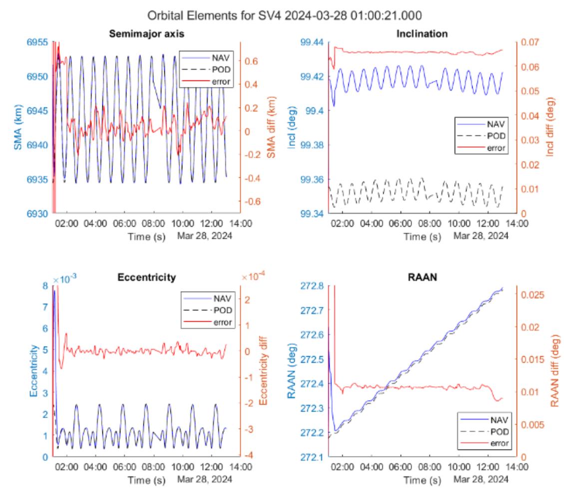

Figure 4 Comparison between onboard navigation solution and flight dynamics precision orbit determination solution.

Figure 4 shows the offset between the onboard navigation filter orbit estimate and the flight dynamics team's precision orbit determination estimate. This indicates that the onboard nav filter has an error in the reference frame conversion from Earth-centered Earth-fixed to the International Celestial Reference Frame. The semi-major axis and eccentricity show close agreement with the flight dynamics orbit solution while inclination and right-ascension show a non-zero offset. Figure 5 shows the long-term trends of the inclination and right-ascension offsets—both have a sinusoidal pattern with a period of 300 days.

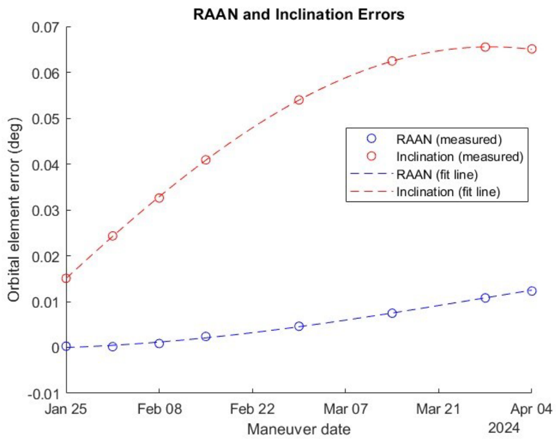

Figure 5 Inclination and RAAN divergence over time

Attempts to correct this offset produced by the filter were unsuccessful, and ROMEO switched to the BESTXYZ message from the Novatel GPS receiver for navigation. Maneuver plans were generated after this switch that met all criteria for execution. However, during test runs of the onboard maneuver plan execution process, with payload commanding of the bus disabled, it was discovered that commands generated by the onboard software were not being sent at the correct time.

During the Starling 1.5 mission extension, discussed further under the Future Work section of this paper, troubleshooting efforts of the navigation filter and maneuver execution timing will continue.

DISTRIBUTED SPACECRAFT AUTONOMY (DSA)

Technology

The collaboration between DSA (its own NASA project funded by the Game Changing Development (GCD) program) and Starling is helping advance autonomy for distributed space systems. Autonomy plays a crucial role in multi-spacecraft missions, enabling spacecraft to make decisions independently, rather than relying solely on ground control. This capability is particularly important for future deep-space missions involving multiple spacecraft, where the communication delays and limited data transmission capacity make traditional command and control approaches impractical.

The Starling flight demonstration component of DSA centers around a GPS Channel Selection Experiment, described in detail in prior papers.[6] This experiment utilizes a dual-band GPS receiver to measure the total electron content (TEC) of the plasma between the spacecraft and GPS satellites. By analyzing these measurements, the experiment aims to capture various phenomena in the ionosphere, such as the Equatorial Ionization Anomaly and the Polar Patches. The DSA system leverages emergent capabilities, namely "shared sampling" and "simultaneous sampling" to optimize the selection of GPS channels across the spacecraft swarm. These approaches enable the efficient allocation of channels for explorative and exploitative observations, maximizing the scientific value of the collected data.

In the flight demonstration, each spacecraft within the Starling system downlinked the complete GPS dataset, regardless of the channel selection. The ground data system then analyzed the data, comparing the actual channel allocations made by the DSA algorithms against an optimal partition. The performance of the algorithms was evaluated based on the degree of match with the optimal channel allocations and the speed at which the DSA system reconfigures in response to changes in the observed features. This experiment was selected as the primary demonstration of the DSA project due to its ability to showcase autonomous reconfiguration in response to natural phenomena without significant integration efforts or modifications to the spacecraft hardware.

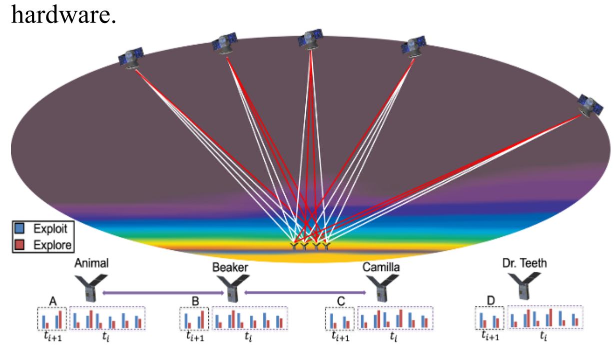

Figure 6 DSA Software performing autonomous GPS channel selection for TEC calculation. A time series below spacecraft A, B, C, and D represents the history of explore and exploit values used by the DSA software to autonomously select GPS channels; red and white lines of sight represented the variation in selected channels.

Figure 6 provides a simplified representation of a channel assignment scenario within the swarm, where multiple spacecraft receive signals from GPS satellites. The experiment involves constraining the number of channels each spacecraft can observe, requiring the DSA system to coordinate channel assignments across the swarm using shared sampling. In the case of spatially constrained phenomena, simultaneous sampling allows multiple spacecraft to observe the features of interest from different vantage points. The performance of the algorithms used in the experiment will be evaluated based on their ability to match the optimal channel allocations and their responsiveness to changes in observed features.

The DSA Flight Software utilizes the core Flight System (cFS) as the framework for each satellite's flight software. This choice ensures compatibility with the Starling payload flight software. For communication middleware, RTI's Connext DDS Micro was chosen due to its ability to handle packet-based communication, scalability, and established flight heritage. The DSA flight mission software consists of three apps within the cFS framework: the Comm App, TEC App, and Autonomy App. The Comm App acts as a wrapper for RTI's Connext DDS Micro, enabling message routing over the Starling MANET; the TEC App processes GPS receiver data and provides inputs to the Autonomy App, which generates a plan for monitoring GPS channels based on inputs from the local TEC App and other satellites.

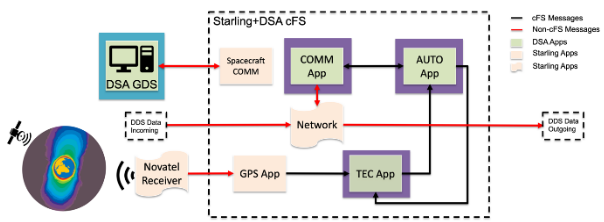

Figure 7 A simplified diagram of the 3 DSA applications operating within the Starling flight software environment, receiving GPS data, and communicating with the DSA Ground Data System (GDS), and utilizing the DDS network.

The flight dataflow diagram (see Figure 7) illustrates the flow of data from raw GPS instrument data to channel selections through the three cFS apps. The Comm App facilitates communication between the local autonomy software and other spacecraft, while the TEC App calculates relevant information from raw GPS range data. The Autonomy App utilizes a Mixed Integer Linear Programming (MILP) solver to find optimal channel allocations by combining rewards from the TEC App and other spacecraft. Additional details regarding these application implementations can be found in prior papers published by the DSA group.[6,7]

The MILP solver was selected to produce channel allocations from a set of satellites in view. MILP offers scalability, fault response capabilities, and seamless integration with the flight software. Other options considered include Hierarchical Task Networks, Temporal Constraint Networks, Markov Decision Processes (Fully and Partially Observable), and Constraint Programming. However, MILP exhibited the necessary features, making it the preferred choice for the DSA experiment.

Flight Results

DSA and Starling consistently established a 3-spacecraft crosslink network via the Comm application, as described in detail within Figure 8. Difficulties in updating the Starling-1 payload software resulted in no 4-node swarm network being achieved during the DSA experiments. The DDS network shares spacecraft state information and group messages. In Figure 8 the gray areas represent time periods with no available data; darker gray represents times where data is missing for more than one spacecraft. (Continuous data does exist onboard the spacecraft, but it is not downlinked sequentially; often the operations team must wait long periods of time before full data is available from Experiment Periods (EPs).) Additionally, the green vertical striped lines represent times when crosslink is turned on and the red vertical striped lines represent times when crosslink is turned off. No data was available from SV1 for this graphic.

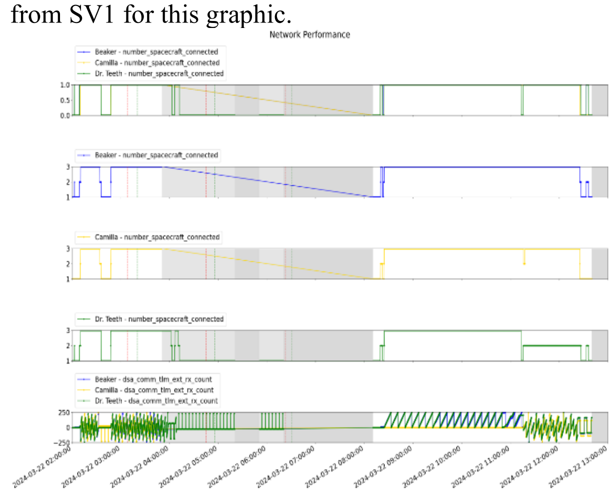

Figure 8 From top to bottom, over a select time period of EP8.1 day 1: Times when all spacecraft are connected to the network = 1.0 and when a full network is not achieved = 0.0; Number of spacecraft connected to Beaker (Starling-2), Camilla (Starling-3) and Dr. Teeth (Starling-4); Accumulated telemetry packets from cross-link on Beaker, Camilla, and Dr. Teeth.

Starling-1 never received a successful DSA software update and was thus never able to join the crosslink network for DSA experiment operations, due to a hardware anomaly that imposed radio uplink limitations. Despite this, DSA's DDS network layer was extremely successful and was the earliest fully successful system that DSA demonstrated.

Critical to understanding the satisfaction of DSA full success criteria is how the project defines "consensus." Consensus exists when all of the spacecraft have the same plan. Periods of non-consensus are expected and can occur due to changes in network topology and configuration, GPS satellite visibility set changes, changes in explore/exploit rewards or changes in the Autonomy app configuration. Changes often occur within a time step that will cause non-consensus, as simulated in prior work.[7]

Consider the interval [01:14:57 - 02:50:38] in Figure 8 above. Consensus is achieved frequently, 893 of 2709 ticks (1 tick = 1 second = a 1Hz cycle of the Autonomy app) in this period but often holds for only a few seconds. Figure 9 demonstrates time until re-consensus for all 3 operational spacecraft. This graphic demonstrates that DSA's autonomy system is "stable," as it tends to regain consensus often as it runs. The mean time until re-consensus over this time was only 2.04 seconds – much better than both the stated goal of 5 minutes and the threshold of 22 minutes.

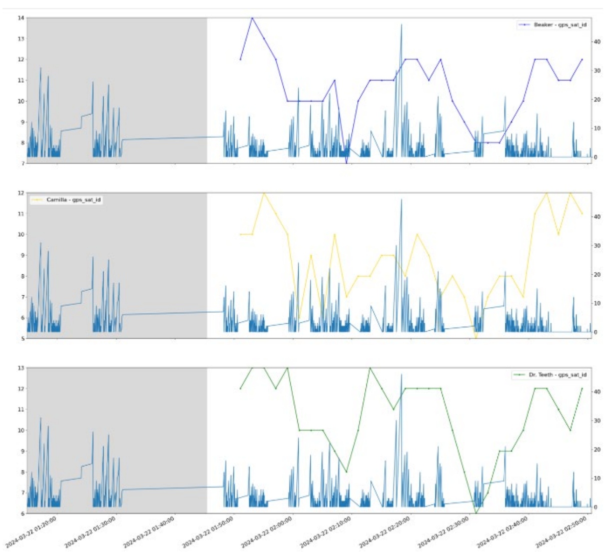

Figure 9: Low frequency plots (blue, yellow, green) represent the total visible GPS satellites per spacecraft. The spiky blue line graph represents the total amount of connective time, measured in seconds, spent in non-consensus.

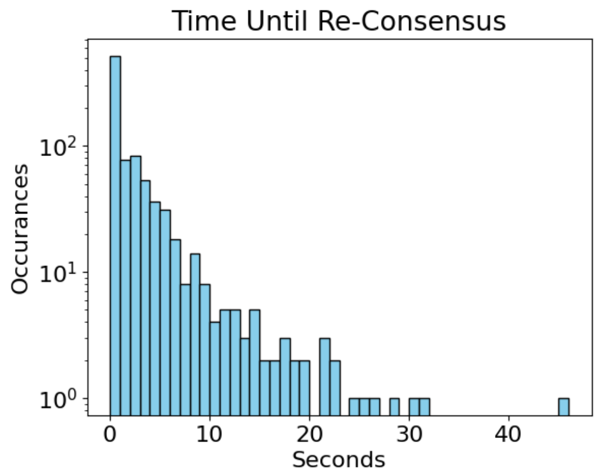

Figure 10 graph demonstrates that, over the selected period of 01:14:57 - 02:50:38 in EP8.1, re-consensus happens quickly (notice the log scale) with a worst-case outlier of 46 seconds

Based on these and other results, DSA can claim to have achieved the first fully distributed autonomous operation of multiple spacecraft and the first demonstration of fully distributed reactive operations onboard multiple spacecraft.

See DSA's 2024 4S paper[8] and 2024 Small Sat paper for additional technical details and measurements of success.

STARLING FORMATION-FLYING OPTICAL EXPERIMENT (STARFOX)

Technology

The software payload for StarFOX is the Absolute and Relative Trajectory Measurement System (ARTMS) developed by the Stanford Space Rendezvous Laboratory (SLAB). ARTMS enables swarms of cooperative observers equipped with low-cost optical sensors to simultaneously estimate their own orbits and the relative orbits of nearby non-cooperative resident space objects using only bearing angle measurements. Bearing angles are obtained by vision-based sensors (VBS) carried onboard the observer spacecraft and may be shared between multiple observers over an inter-satellite link (ISL). GPS measurements may optionally be applied by observers to determine their own absolute orbits with higher accuracy and timeliness. An example ARTMS scenario is presented in Figure 11.

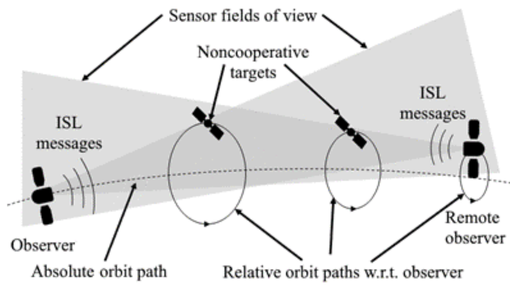

Figure 11: An example scenario featuring two cooperative observers and two passive targets

Angles-only navigation in orbit is typically characterized by weak observability due to the lack of range information provided by the measurements. To overcome this, ARTMS applies three core algorithms. First, Image Processing (IMP) leverages domain-specific modeling of target kinematics and multi-hypothesis techniques to identify and track multiple targets in images from a single camera without a-priori relative orbit knowledge. Second, a new Batch Orbit Determination (BOD) algorithm applies a relative orbital element state definition to provide initial estimates of the orbits of all spacecraft in the local swarm, using a one-dimensional sampling scheme to resolve the weakly-observable range to each target. Third, a Sequential Orbit Determination (SOD) algorithm based on the unscented Kalman filter continuously estimates the orbits and auxiliary parameters of the local swarm by seamlessly fusing measurements from multiple observers shared over an ISL. The ARTMS instance on each observer can autonomously 1) form batches of bearing angles corresponding to each visible target, 2) initialize an orbit estimate for each target, and 3) refine the complete swarm state estimate over time. The only *a priori* information required is a coarse absolute orbit estimate for each observer. A high-level block diagram of the three core ARTMS algorithms and their data flow is presented in Figure 12.

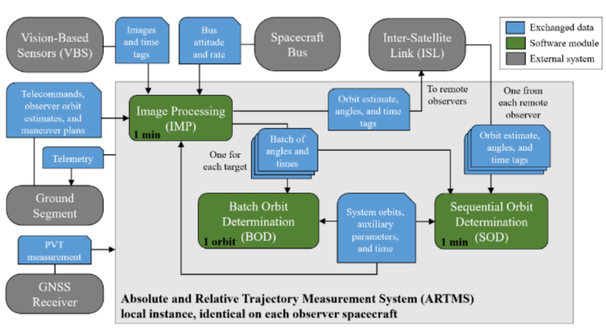

Figure 12: ARTMS angles-only navigation architecture block diagram

ARTMS has been designed to impose minimal hardware and software requirements on a system. For example, the StarFOX concept of operations specifically applies star trackers as the VBS, which are a fundamental attitude determination and control system (ADCS) sensor on nearly all modern satellites. Thus, no additional hardware is needed, only the ability to run a star tracker in "image capture" mode, which was an inherent capability of the BCT XB1 6U bus used by Starling. The ARTMS algorithms are implemented in software as three distinct cFS apps which execute at fixed sample times on all spacecraft throughout the experiment period. During StarFOX experiments in LEO, image measurements were obtained and processed every 1-2 minutes and ISL transmissions occurred every 1-2 minutes (when active). The BOD app executed once per orbit.

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

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

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

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

Mode 4: Multi-observer, fully autonomous: multiple observers share measurements over the ISL to estimate target relative orbits with improved accuracy, and to estimate the absolute orbit of each observer (i.e. no GPS usage)

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

Flight Results

All four navigation modes were demonstrated during flight operations. For the sake of brevity, a subset of the overall results is described here. A representative flight image (as applied for navigation) is illustrated in Figure 13, with point sources classified by the IMP algorithm.

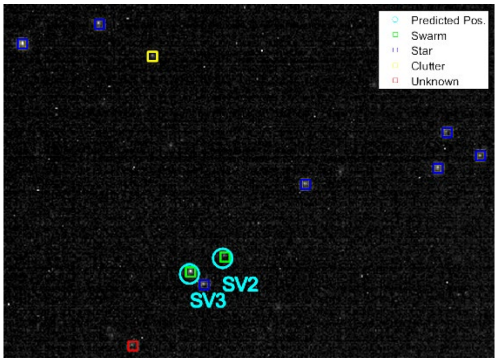

Figure 13: An in-flight star tracker image with two swarm targets visible

Navigation Mode 2 was demonstrated on 04/09/24 by Space Vehicle 4 (SV4) tracking SV3, using in-flight measurements processed by ARTMS on the ground. Recall that in this mode, no *a priori* knowledge is provided regarding the relative orbits of targets. Instead, the observer spacecraft autonomously begins tracking target candidates using IMP. After forming a batch of measurements (over several orbits), it computes an initial relative orbit estimate using BOD. This estimate is used to initialize the SOD algorithm and is later refined as more images arrive. The SV4-SV3 intersatellite range during this period was approximately 80 km.

Figure 14 provides an example of a target range estimate produced by BOD. The initial range estimate, computed using 37 target measurements, possessed an error of approximately 2 km and an uncertainty of 18 km. Later estimates incorporating larger measurement batches typically display improved uncertainties. Pre-flight requirements were that BOD should produce initial range errors of less than 20% of the true target range,[9] which is fulfilled by the in-flight result. This is the first demonstration of using flight data and flight software to autonomously initialize an angles-only navigation system in orbit.

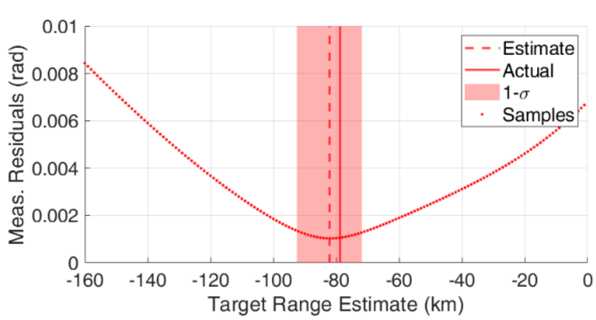

Figure 14: An initial target range estimate produced by BOD

Figure 15 provides an example of target state estimation performance after an initialization by BOD. Errors and uncertainties are presented in the RTN frame of SV4. Bearing angles provide comparatively weaker observability of target range, which typically leads to larger along-track uncertainties. In addition, at longer ranges, the same bearing angle measurement noise corresponds to a larger geometric uncertainty, which implies proportionality between range and the overall state uncertainty. In Figure 15, uncertainties are primarily distributed in the along-track direction as expected, with final relative state uncertainties (1σ) of ~1% of target range. This similarly meets pre-flight performance requirements.[9] Note that visibility issues on orbit (caused by dimmer-than-expected targets and higher-than-expected image noise) resulted in long measurement outage periods, as indicated by the white vertical regions. Nevertheless, ARTMS was able to produce reasonable state estimates, without requiring maneuvers to disambiguate target range.

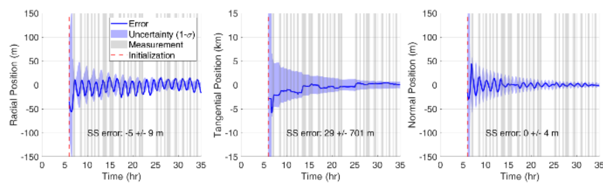

Figure 15: Target relative state estimation errors and uncertainties after BOD initialization

Demonstrations of Navigation Mode 4 were attempted on 05/09/24 and 05/16/24 in flight but hardware issues prevented successful execution. It was later demonstrated using flight data from 03/24/24, post-processed by an ARTMS flight software 'digital twin' running on the ground. The scenario applied three ARTMS instances to process approximately 3300 star tracker images, with simulated crosslink measurements between them. The on-board state estimate for each simulated spacecraft was initialized using artificially degraded orbit knowledge, to emulate a coarse initialization (with initial uncertainties on the scale of several kilometers). After initialization, no external absolute orbit knowledge is provided to any spacecraft (i.e. no GPS data is processed).

Absolute orbit estimation using only inter-satellite bearing angles is particularly challenging due to the weak observability of the problem. However, if multiple observers are present and are obtaining measurements of a common target (and each other), enough geometric is present to facilitate long-term absolute orbit estimation without external updates.[9] Figure 16 presents absolute position errors and uncertainties in RTN, estimated by SV4. SV4 observes SV2 and SV3 as targets, but also receives crosslink measurements from SV3 (measuring SV2 and SV4) and SV2 (measuring SV3). Initial absolute position uncertainties of [3, 6, 3] km in the RTN coordinate frame are consistently reduced throughout the simulation, with final uncertainties after 36 hours of approximately [1, 2, 1.6] km. Errors also remain within 1σ uncertainty bounds, indicating reasonable filter health. If initial orbit errors are less degraded with, e.g. position uncertainties of [1, 2, 1] km in RTN, final position uncertainties in the on-ground experiment were approximately 850 m.

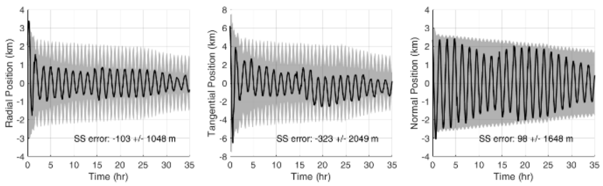

Figure 16: Multi-observer absolute orbit estimation performance for SV4

Figure 17 presents the epochs at which measurements between a specific spacecraft pair were recorded. The most consistently visible link was SV4 observing SV2, whereas other measurement combinations were sparser (due to the visibility issues previously described).

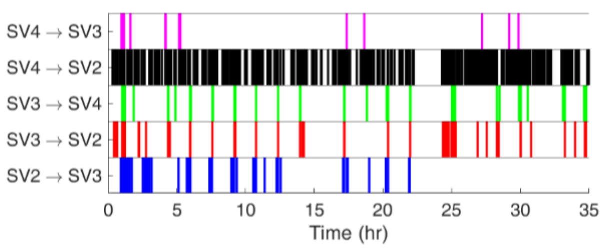

Figure 17. Measurements produced by the swarm during a multi-observer case

Use of multiple cooperative observers also provides advantages for relative navigation, because the additional geometric information significantly reduces target range ambiguity. For similar swarm geometries in successive experiments 02/05/24 and 02/11/25, multi-observer navigation was able to reduce relative state estimation uncertainty by more than a factor of two, compared to the single-observer case.

Overall, StarFOX has successfully conducted the first in-flight demonstrations of autonomous angles-only navigation for a spacecraft swarm, including multi-target and multi-observer swarm scenarios; autonomous initialization of relative orbit estimates in flight; and angles-only absolute orbit determination.

FUTURE WORK

The Starling primary mission concluded in May 2024, having achieved most of its primary objectives. The team is now embarking on a mission extension called Starling 1.5, the designation referring to an implementation enabled solely through ground and onboard software updates. The following are some examples of the added scope with the Starling 1.5 mission extension:

1) Demonstrate a Space Traffic Coordination architecture in cooperation with the SpaceX Starlink constellation. Over the past few years, Low Earth Orbit has become increasingly crowded, and furthermore spacecraft are becoming increasingly capable with autonomous maneuvering capabilities. Previous methods for detecting conjunctions between spacecraft through the $18^{th}/19^{th}$ Space Defense Squadrons and manually coordinating maneuvers between different owner/operators is becoming increasingly untenable. To address this issue, an architecture has been developed where the spacecraft and supporting ground systems can themselves negotiate maneuver responsibilities through a ground hub to avoid potential collisions. Together with the SpaceX Starlink satellites, the Starling spacecraft will demonstrate the viability of the system.
2) Continue advancement of Distributed Spacecraft Autonomy developed during the primary mission to demonstrate “send and forget” autonomy. This will be accomplished by adding autonomous control of the spacecraft hardware such as the crosslink radios to the capabilities of the onboard system. In addition, the science data will be refined via onboard edge computing to more precisely detect short-lived and transient electron density “bubbles” in the upper ionosphere, turning on and off the power consuming crosslink radios in response to these science observations.
3) Extend the capabilities of the StarFOX onboard navigation experiment including a) detection and tracking of unknown satellites for space domain awareness; b) faster initialization for quickly evolving scenarios such as space rendezvous; c) estimation of auxiliary state components such as clock errors, drag coefficients and solar radiation pressure coefficients; and d) achieve higher accuracy via sensor fusion of optical and crosslink radio ranging measurements, enabled via a firmware update to the crosslink radios.
4) Extend the MANET ad hoc network with capabilities that will be needed for deep space, including Delay Tolerant Networking, and a Store and Forward overlay for the CCSDS based file transfer protocol.

CONCLUSION

The Starling technology demonstration operated from July 2023-May 2024 and demonstrated key enabling technologies for future swarm missions. The project completed most of its key objectives, achieving full success on three out of five mission objectives and minimum success on the other two.

The Better Approach To Mobile Ad-hoc Networking (B.A.T.M.A.N.) protocol was implemented within the onboard software and used to demonstrate the discovery and establishment of a network without the spacecraft having *a priori* knowledge of their neighbors. The network was used to demonstrate point-to-point communications, data routing of messages when point-to-point communication was unavailable, and file transfers between spacecraft. The network was characterized using multiple modulation schemes and data rates.

In terms of managing the spacecraft swarm formation, manually commanded maneuvers from the ground were successfully used to bring the swarm into both the in train and Passive Safety Ellipse formations, and to maintain inter-spacecraft distances of 50-100 km. However, the goal was to demonstrate formation change and station-keeping maneuvers through the Reconfiguration and Orbit Maintenance Experiments Onboard (ROMEO) onboard CFA software. Onboard maneuver plans were generated, but these could not be executed due to issues with the planned propellant usage and timing. This capability will continue to be investigated during the Starling 1.5 mission extension.

The Distributed Spacecraft Autonomy (DSA) experiment demonstrated autonomous collaboration and consensus between three spacecraft swarm members in conducting science observations of the Earth's ionosphere. The DSA experiment is the first demonstration of a fully distributed onboard reasoning system, able to react quickly to changes in science observations. Unfortunately, the demonstration could not be accomplished across all four spacecraft due to space-to-ground communication issues on the one spacecraft previously mentioned, and limitations on the size of files that could be side loaded across the network for the requisite software update.

The Starling Optical Formation Flying Experiment (StarFOX) demonstrated both relative and absolute orbit determination by using optical data from onboard star trackers. The software was able to recognize and track fellow swarm members in the presence of other spacecraft, debris, and background stars. Using angles-only algorithms, the swarm was able to determine each member's orbit and could improve accuracy by sharing navigation data between members over the ad hoc network. StarFOX is the first-ever published demonstration of this type of swarm navigation system.

The Starling mission represents a significant milestone in the development of spacecraft swarms. Spacecraft swarms have the potential to revolutionize space science and exploration for future NASA missions. In fact, Starling is considered a technology precursor mission for the Helioswarm project which will be deploying a swarm of spacecraft to study solar plasma turbulence in 2029.

Finally, the Starling mission is also a tribute to the large leaps in technology that can be gained by utilizing small, low-cost spacecraft and taking appropriate risks. Through the Starling mission and its follow-on activities, these four small CubeSats will have a significant impact on NASA and the space industry.

Acknowledgments

Starling is funded by the Small Spacecraft Technology (SST) program within the NASA Space Technology Mission Directorate (STMD), with funding for Distributed Spacecraft Autonomy (DSA) provided by Game Changing Development (GCD) within STMD. The team would also like to acknowledge the following partners vital to making the mission a success: Blue Canyon Technologies, Rocket Lab USA, Stanford University Space Rendezvous Lab, York Space Systems (formerly Emergent Space Technologies), CesiumAstro, and L3Harris Technologies, Inc.

References

1. Hanson, J., Chartres, J., Sanchez, H., Oyadomari, K.; "The EDSN Intersatellite Communications Architecture"; Proceedings of the 28th Annual AIAA/USU Conference on Small Satellites, 2014, paper: SSC14-WK-2.
2. NASA, "Nodes - Network & Operation Demonstration Satellite," 6 August 2017. [Online]. URL: https://www.nasa.gov/centers/ames/engineering/projects/nodes.html, [Accessed 31 May 2018].
3. Sanchez, H., McIntosh, D., Cannon, H., Pires, C., Sullivan, J., D'Amico, S., and B. O'Conner, "Starling1: Swarm Technology Demonstration," 32nd Annual AIAA/USU Conference on Small Satellites, 2018.
4. D'Amico, S. and O. Montenbruck.; "Proximity Operations of Formation-Flying Spacecraft Using an Eccentricity/Inclination Vector Separation," Journal of Guidance, Control, and Dynamics. Vol. 29, No. 3, May-June 2006.

5. Emergent, “Cluster Flight Application Suite,” [Online]. URL: https://www.emergentspace.com/products/cluster-flight-application/. [Accessed 3 June 2024].

6. Cellucci, D., Cramer N.B., and Jeremy Frank; “Distributed Spacecraft Autonomy,” ASCEND 2020 paper: AIAA 2020-4232.

7. Cramer, N., Cellucci, D., Adams, C., Sweet, A., Hejase, M., Frank, J., Levinson, R., Gridnev, S., and Lara Brown, “Design and Testing of Autonomous Distributed Space Systems,” In Proceedings of the AIAA/USU Conference on Small Satellites, 2021.

8. Adams, C., Frank, J., Gridnev, S., Iatauro, M., Kempa, B., Lassiter, C., and Richard Levinson, “Advancing Autonomy in Distributed Space Systems: Insights from On-Orbit Testing with the Starling 1.0 Mission,” In Proceedings of the Small Satellites Systems and Services (4S) Symposium, 2024.

9. Kruger, J., Koenig, A. W., and D’Amico, S.; “Starling Formation-Flying Optical Experiment (StarFOX): System Design and Preflight Verification”, Journal of Spacecraft and Rockets, vol. 60, no. 6 (2023).