SSC24-VII-01

Flight Results and Lessons Learned from the Starling Propulsion System

Terry Stevenson
KBR, Inc., NASA Ames Research Center
Moffett Field, CA 95054; 650-604-6321
terry.h.stevenson@nasa.gov

ABSTRACT

NASA's Starling mission is a swarm of four 6U Cubesats flying in formation in Low Earth Orbit to demonstrate scalable swarm technologies, particularly related to crosslink networking. In order to maintain and modify their formation, the Starling spacecraft use a cold gas propulsion system called Hamlet, which provides each spacecraft with 30 m/s of delta-V, as well as attitude control during maneuvers. Hamlet is a fully self-contained system, incorporating propellant tanks, valves, nozzles, and control electronics. Most of the structure of Hamlet is a single piece of stereolithography-printed composite, which simplifies assembly and allows for unusual tank geometry that maximizes propellant volume in the allocated space. An extensive qualification campaign was conducted for Hamlet, including performance characterization that has been largely validated by in-flight measurements. Since the start of nominal mission operations in late summer 2023, Hamlet has conducted weekly or bi-weekly maneuvers of up to 0.29 m/s to assemble the swarm, maintain formation, and change swarm configurations as required by the experiments. Operations have not been without challenges, including a propellant leak on one unit and thrust variability issues on all four. This paper describes the design and implementation of Hamlet, as well as in-flight performance data, anomaly investigation and resolution, and lessons learned.

INTRODUCTION

Starling is a four-spacecraft mission run by NASA Ames Research Center to advance the state of the art in multi-spacecraft swarms, with a particular focus on technologies that involve intersatellite communication to coordinate actions across the swarm.

Starling uses the 6U XB1 Cubesat bus, built by Blue Canyon Technologies. The crosslink radios and antennas were provided by CesiumAstro. The payload processor, power supply, and payload flight software were developed in-house at NASA Ames. NASA Ames also developed the cold gas propulsion system, called Hamlet. System-level integration and testing was performed at Ames before delivery to Rocket Lab for launch.

Starling Mission Experiments

There are four separate experiments on Starling, all related to advancing the ability of small spacecraft to cooperatively and autonomously operate in formation.[1]

The MANET (Mobile Ad-hoc NETwork) experiment demonstrated a flexible network that could automatically change network topology in response to changes in formation and changes in spacecraft state (for example, routing around a spacecraft in safe mode). This network formed a key support layer for the other three experiments.

The StarFOX (Starling Formation-Flying Optical Experiment) experiment, run in partnership with Stanford University, demonstrated angles-only relative navigation using images taken by the XB1's star trackers.[2] The experiment was able to identify and track the other Starling spacecraft using these images, and to correctly estimate their trajectories.

The ROMEO (Reconfiguration and Orbit Maintenance Experiments Onboard) experiment attempted to demonstrate autonomous maneuver planning and execution. ROMEO was able to generate maneuver plans, but these could not be executed due to excessively high delta-V and timing issues. Improvements are ongoing for the Starling 1.5 mission extension.

The DSA (Distributed Spacecraft Autonomy) experiment demonstrated cooperative experiment execution by the Starling swarm to measure total electron count in the ionosphere via GPS measurements.[3] The spacecraft communicated via crosslink radio to coordinate measurements and identify times and locations of interest, and the framework for this experiment is applicable to other types of measurements requiring coordination between spacecraft with minimal ground supervision.

Mission Propulsion Need

All of these experiments require the Starling spacecraft to be in a relatively close formation, both to allow crosslink communication with low-gain antennas and to allow visual tracking for StarFOX. The initial formation is a string of pearls formation, where the four spacecraft follow the same orbital path, spaced out in the in-track direction. Later in the mission, the swarm was moved to a passive safety ellipse (PSE) formation, described below.

A propulsion system was required to maintain this formation with appropriate spacing, demonstrate autonomous maneuvering, and reduce the orbital decay rate. After a survey of available systems, the project decided to design and build one in-house, drawing on experience from developing the BioSentinel propulsion system.

Given the relatively low thrust and delta-V requirements for the mission, a cold gas system was chosen for its lower cost and lower risk to other payloads aboard the rideshare.

BioSentinel Heritage

The Hamlet propulsion system was based on a previous design flown on the BioSentinel mission. BioSentinel was also a 6U Cubesat, launched as a secondary payload on Artemis 1 in November 2022 to study the interplanetary radiation environment.

A seven-nozzle cold gas system was flown on BioSentinel to provide attitude control for detumble and reaction wheel desaturation. A limited delta-V capability was also present for a potential lunar avoidance maneuver, but this was ultimately not required and the delta-V capability remains untested. As of June 2024, BioSentinel remains operational, with the propulsion system used roughly every four weeks for momentum unloading.[4]

PROPULSION SYSTEM DESIGN

Hamlet is approximately 2U and 1.47 kg dry (2.45 kg fully loaded). It can provide 30 m/s of delta-V from four nozzles and can provide 3-DOF attitude control during maneuvers. The system is built around a large, hollow 3D-printed structure, with metal manifolds and attached control electronics.

Hamlet is designed as a burn-refill system with two propellant tanks. The larger (799 cm^3) tank, called the main tank, holds the bulk of the propellant as a saturated liquid-vapor mixture. Two solenoid valves allow propellant to flow from the main tank to the smaller (41 cm^3) tank, called the plenum. The plenum is maintained roughly 10°C warmer than the main tank, ensuring that it remains entirely vapor rather than liquid-vapor mixture. The four nozzles are then fed by individual solenoid valves from the plenum.

This two-tank arrangement with nozzle valves being fed from the vapor-only tank prevents the nozzles from ingesting liquid droplets, which can lead to significant variability in performance. Figure 1, below, shows a process diagram of Hamlet.

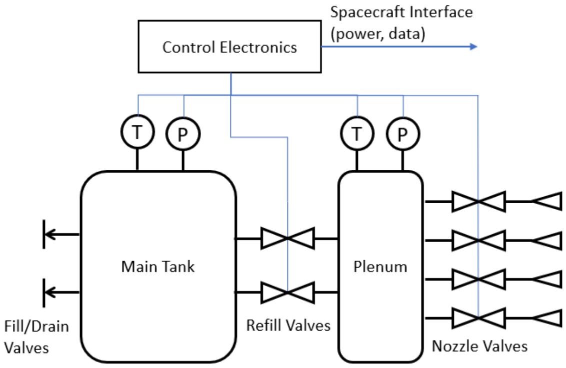

Figure 1: Hamlet P&ID diagram, showing two propellant tanks, valves, sensors, and nozzles.

The four nozzles are pointed in the Z+ direction (Z- thrust direction) angled 20° inward towards the X-Z plane (this can be best seen in Figure 3). This angling gives Hamlet the ability to produce 3-DOF torques during maneuvers via off-pulsing and allows ACS to control attitude entirely with the propulsion system. This does, however, sacrifice one-out valve redundancy since all four are needed for zero net torque.

SV4’s assembled propulsion system is shown in Figure 2. The two long bayonet valves on the left of the image are the fill/drain ports. The upper port has a probe installed to equalize the tank with the surrounding atmosphere. These probes were installed whenever the tanks were unloaded to serve as a visual reference for lab personnel, showing that a tank was not pressurized and could be handled without additional PPE. Sensor cabling and electronics shielding were disconnected before assembly to allow staking to the bus structure.

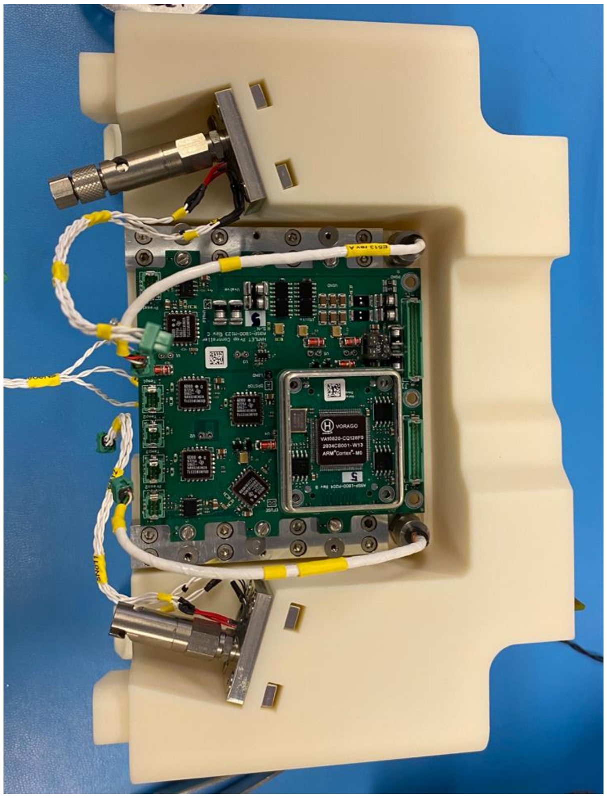

Figure 2: SV4 Hamlet prior to system integration. Electronics EMI shielding is not yet installed.

Propellant

Like BioSentinel, Hamlet uses R-236fa as propellant. R-236fa (1,1,1,3,3,3 hexafluoropropane) is a non-toxic, non-flammable refrigerant with a relatively low saturation pressure of 583 kPa at 50°C.[5] The low storage pressure and lack of chemical hazards simplifies propellant handling during spacecraft integration as well as secondary payload safety approvals. The non-toxic nature of the propellant also allows limited use of the prop system in a lab environment without PPE beyond oxygen monitoring. This allows prop system-in-the-loop testing and final pre-delivery verification of free flow through the system.

R-236fa performs somewhat poorly on a mass basis, with a specific impulse of only 42s. However, its high density (1.27 g/cm^3 at 50°C) allows a larger quantity of propellant to be stored, giving a volume-basis performance better than other common cold gas propellants such as butane. This is particularly important for highly volume-constrained designs, which is common among Cubesats.

Propellant Tanks

Hamlet's propellant tanks, pipes, nozzles, and manifold gasket grooves are all contained in a single piece of Somos PerForm, a high temperature stereolithography-printed particle composite. The use of 3D printing for the tanks allows more flexibility in the tank's internal geometry as well as nozzle placement, giving greater propellant volume for the same overall envelope. However, the material is weaker than aluminum or other more traditional pressure vessel materials, so the system is limited to lower pressure propellants.

One of the printed structures with no attached hardware is shown in Figure 3 below. Note the channel and "mousehole" feature on the bottom center of the structure. These are cutouts of the allocated prop volume to accommodate a solar panel release mechanism and antenna RF cable, respectively. Accommodating such features is relatively simple with a printed propellant tank.

Figure 3: Hamlet printed structure. Note the inward angling of the four nozzles at the corners of the visible face.

This approach was successfully used in the BioSentinel propulsion system, where it was necessary to fit all of the nozzles, piping, and propellant storage into a very restricted volume. Hamlet required more propellant than the BioSentinel system and had a larger volume allocation, but still benefited from the complex internal geometry of a printed structure. This flexible geometry also allowed Hamlet to more efficiently use its allocated volume, which was deformed around three patch antennas and their associated cabling, as well as the solar array release mechanism.

A negative-space CAD view of the printed structure is shown in Figure 4, showing the tank volumes and fluid routing. In addition to allowing more flexible routing, incorporating the pipes into the printed structure also reduces the number of pressure seals that must be made, which simplifies assembly and reduces the number of possible leak locations.

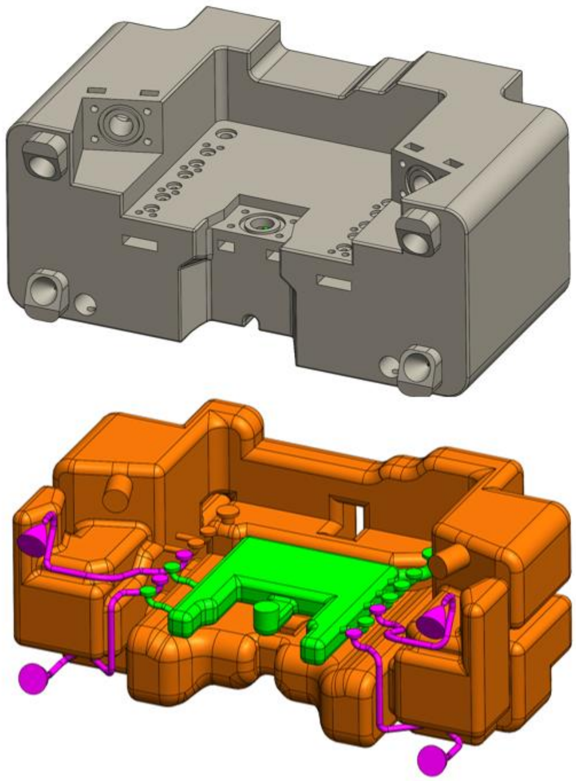

Figure 4: Hamlet printed structure (top), and negative-space view from the same angle (bottom). Main tank and connected piping are in orange, plenum in green, and nozzles in magenta.

Both tanks were designed to a 2.5X maximum pressure standard (225 psi), with the condition that maximum triaxial stress must not exceed measured tensile strength. The triaxial condition was enforced because the expected failure mode of the printed structure is brittle fracture, given its low elongation to break of 1.2%.[6]

A total of 16 structures were printed, 14 of which were proof tested to 135 psi (1.5X maximum expected pressure). The other two printed structures were rejected before proof testing, one for a small pinhole leak and one for a partially obstructed nozzle pipe. CT scanning of the printed parts was used to verify nozzle geometry and lack of pipe obstructions.

Manifolds and Valves

Four aluminum manifolds provide interfaces into the propellant tanks for sensors and valves. These manifolds are face-sealed to the printed structure with Buna-N gaskets seated in printed grooves on the structure. Square profile O-rings were used for all metal-to-printed seals due to the higher surface roughness of the printed part.

Three of the manifolds cover single tank openings: two main tank openings for heaters and fill/drain valves, and one plenum opening for heaters.

The fourth manifold is more complex, and contains the six solenoid valves and two pressure sensors, with 14 interfaces into tanks as well as all four nozzle pipes. The control electronics are mounted to this manifold to minimize path length between the valves and their driving circuits. A CAD view of this manifold is shown in Figure 5, note that the tank interfaces are on the bottom of the manifold and not visible in this view. See Figure 4 for overall locations of the manifolds.

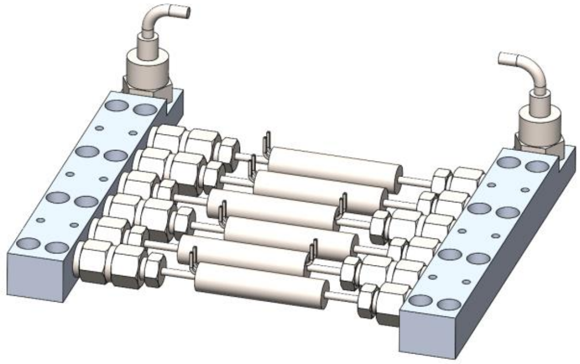

Figure 5: CAD image of Hamlet valve manifold, with pressure sensors at back.

Each of the six solenoid valves has a 5µm sintered filter immediately upstream to prevent particles from lodging in the valves and causing them to fail open. Although the tanks were cleaned as well as possible with IPA flushing, the complex internal geometry prevents direct inspection and verification of cleanliness. Debris left over from the printing process and bead blasting is almost certainly still present in small amounts in all of the flight tanks. The valve manifold components, by contrast, are all traditionally manufactured pressure fittings and can be precision cleaned with normal methods. The fluid paths within Hamlet are separated into cleanable and non-cleanable areas, with the valves protected by filtration from the non-cleanable areas (printed piping and tank interiors). This approach has been successful on both BioSentinel and Hamlet, with no debris-related valve sticking.

Controller

The control electronics are installed directly onto the valve manifold. These boards contain the valve driving circuits, sensor ADCs, and a Vorago VA10820 processor. The VA10820 is far more powerful than required for this application but represented another technology demonstration opportunity for a rad-tolerance Cubesat-scale processor and fulfilled the radiation tolerance requirement for the prop system processor.

Hamlet communicates with the XB1 via serial line, receiving thrust commands and sending propellant and valve state telemetry. A separate GSE port gives a second serial line for ground debugging and reprogramming. Hamlet is not reprogrammable in flight.

The thruster off-pulsing controller is a legacy XB1 controller from a previous Blue Canyon mission, chosen to save development time and cost. The legacy controller does not accept telemetry from the prop system, so maneuvers are commanded blind with no knowledge of the current propellant state. Since thrust is highly dependent on temperature, the bus controller commands Hamlet assuming a minimum operating thrust of 4.2 mN per nozzle. Hamlet's processor then scales these commanded times down based on the current propellant temperature to provide the correct average thrust. This commanding cycle takes place at 1 Hz, which is sufficiently fast to maintain spacecraft pointing to within 0.5° during maneuvers.

The use of this legacy system saved significant development time but is a somewhat awkward fit for Hamlet's temperature-dependent thrust and limits the average thrust achievable during maneuvers.

GROUND TEST CAMPAIGN

All Hamlet flight units went through a full acceptance test campaign before system integration, and a qualification campaign was conducted on a flight-like engineering unit (EDU) before flight unit procurement.

All units (including the EDU and spare tanks) were proof pressure tested per AIAA S-080/S-081, with 50 cycles from 0 to 90 psi (the maximum expected in-flight pressure) and a final cycle to 135 psi with a five minute dwell. Propellant was loaded and the unit was leak checked with a fluorine sniffer. Each unit then went through vibration testing and thermal vacuum (TVAC) testing including fluorine leak sniffing of the vacuum chamber outlet pumps to confirm no propellant loss at the temperature extremes. Each unit was then tested on a torsional pendulum thrust stand, with a minimum of 100 pulses per nozzle to measure thrust and specific impulse.

NASA standards for additively manufactured polymer pressure vessels did not exist during Hamlet development, so the project adopted a conservative approach and required testing of the flight units to GEVS protoflight levels, despite testing of the EDU to GEVS qual levels. This was intended to account for variability in the printing process from unit to unit and was applied to vibration levels and TVAC temperature limits.

After integration, the spacecraft went through vibration testing and TVAC at the system level, so each flight unit experienced these environments twice before launch. Final spacecraft checkout before delivery to Rocket Lab included a brief flow test of each nozzle as well as leak screening with a fluorine sniffer.

MANEUVER CONOPS

Each maneuver is made up of one or more burns, defined here as continuous periods of propulsion. For stationkeeping and orbit raise maneuvers these burns are typically spaced half an orbit apart (for example, to raise and then recircularize the orbit). Crosstrack burns to adjust inclination have different timing requirements, and maneuvers containing both may have burns spaced only minutes apart.

Maneuver Timing

Ninety minutes before the first burn of a maneuver, the spacecraft's Globalstar radio beacon is disabled, ensuring minimal interference with GPS measurements to ease maneuver reconstruction. Hamlet is preheated to its operational temperature of 40°C. While the main tank heaters are on during this period, there is far too much thermal mass in the main tank to effectively heat it in a reasonable time. The focus of this preheat is to warm the plenum and valve manifolds, since these are the strongest drivers of performance. The plenum heater can generally keep the plenum warm once preheated, despite evaporative cooling during refills.

Ten minutes before each burn, a preparation script is run that rapidly cycles all of the nozzle valves. The overall open time is low enough to not impact the trajectory, but the rapid cycling can potentially unstick valves. Five minutes before the burn, the spacecraft is commanded into the burn attitude.

During the burn, Hamlet receives commands from the bus at 1Hz containing valve timings. Hamlet adjusts these commands based on its current temperature, then opens the nozzle valves for a maximum of 500 ms. The different nozzle valves do not necessarily open for the same duration, with some closing early to provide attitude control (off-pulsing). The refill valves are then opened for up to 500 ms to repressurize the plenum, and the cycle repeats.

After the burn is complete, if any further burns remain, Hamlet is kept pre-heated during the waiting period. If all burns are complete, Hamlet is powered off and heaters are returned to their survival setpoints. Globalstar is reenabled ninety minutes after the last burn, ensuring that one complete orbit of GPS measurements free of interference are available before and after each burn.

Maneuver Planning

As with other aspects of Starling mission operations, maneuver planning is done on a weekly cadence in order to accommodate a nominal 5 day work week. Formation goals are discussed early in the week and chosen based on intersatellite drift and the needs of experiments running that week.

In nominal cases, maneuver scripts are generated and reviewed on Wednesday. They are uplinked and loaded over the next 24 hours, with maneuver execution occurring Thursday night or Friday morning.

Prop system telemetry and GPS data are downlinked on Friday morning, and the flight dynamics team uses the GPS data to reconstruct the maneuvers and determine actual versus predicted performance.

A disadvantage of this cadence is that in cases where performance was poor and a second maneuver is needed to maintain appropriate spacing, planning and execution of this maneuver must take place on the weekend.

OPERATIONS

The Starling swarm launched on July 18, 2023 on the 39th Rocket Lab Electron mission along with 3 other spacecraft. Contact was established shortly afterwards with all four spacecraft. SV1 was found to have a radio issue (later discovered to be a pre-launch fault) that greatly shortened its ground contacts and required more time to commission.

Spacecraft 1 Leak

During the commissioning phase, a propellant leak was discovered on SV1. The first indication of this was from momentum telemetry, which showed significantly greater momentum accumulation than the other three vehicles, requiring higher torque rod duty cycling to maintain low reaction wheel speeds. The torque was determined to be fixed in the spacecraft body frame, which ruled out residual magnetic torque and strongly pointed towards a propellant leak.

This was confirmed after GPS data became available. Orbit reconstruction showed a consistent body-fixed unmodeled force disturbing the orbit, with a magnitude of approximately 50 μN, larger than explainable by any other disturbance forces. Assuming a pinhole leak (and thus a near-sonic exit velocity), this results in the loss of 31-34 grams of propellant each day.

The most likely cause of the leak was initially judged to be damage to the printed structure during launch. Such damage could not be repaired, so for anomaly response purposes it was set aside in favor of an O-ring seal failure. Such a failure could potentially be reversed by additional heating of the system.

This was attempted in the first prop system commissioning activity, in which the prop system was powered on for four hours and its associated heaters were set to their higher operational limits. By the end of this period the momentum accumulation was greatly reduced, in family with the other vehicles. The orbit during this period was also reconstructable with no unmodeled forces. The leak returned slowly over the three hours following this activity.

After this confirmation, the prop system was powered back on and its survival heater setpoints were raised to maintain this higher temperature. The leak again disappeared with the higher temperature, and momentum accumulation and orbit residuals remained normal. The prop system electronics were kept powered on to provide waste heat directly to the valve manifold, since the three heaters are all located in the propellant tanks with poor conduction paths to the valves.

On two occasions before commissioning on SV1 was complete, spacecraft resets caused the prop system to be powered off and its heaters reset to pre-launch setpoints. In both cases the prop system was powered back on during the next ground contact, but the leak was observed to reappear during the hours between reboot and ground contact. The spacecraft startup script was modified to automatically power on prop and increase its heater setpoints as part of the boot sequence, and the leak has not been observed since August 2023.

Troubleshooting of this leak was hampered by the SV1 radio anomaly, and the leak was not fully resolved until 16 days after launch. This resulted in an estimated loss of 543 grams, over half of the pre-launch propellant mass.

All four units were screened for leaks immediately before delivery and none were found. Given the magnitude of the leak (1.29 grams per hour, 10,000 times greater than the screening sensitivity) and the fact that propellant remains in SV1, the leak must have started during or shortly after launch and deployment. The resolution of the leak by raising the tank temperature strongly indicates an O-ring seal issue, since higher temperatures can improve elasticity and higher internal pressures can force gaskets into the seal more fully.

The seal failure is believed to be caused by low temperature, as SV1 experienced colder temperatures during the first day of the mission than the other three vehicles. Due to the radio issue the actual minimum temperature experienced by the valve manifold is not known, but is less than -15°C, which was the coldest temperature at which prop was tested on the ground.

The location of the leak has been tentatively identified using the disturbance torque on the vehicle. Given the size of the leak, it must have originated in the main tank as the plenum cannot store sufficient propellant. There are three O-ring face seals between the valve manifold and the main tank that are in the correct location to produce the observed torque. Two of these are at the upstream side of the refill valves and the third is around the main tank pressure sensor. These three locations are shown in Figure 6.

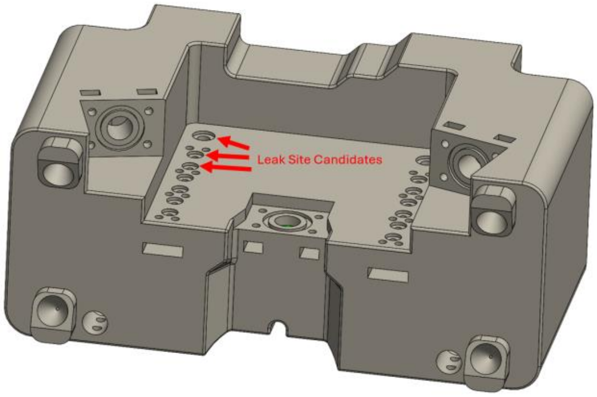

Figure 6: Three possible seal locations identified as the cause of the SV1 propellant leak.

Spacecraft 1 Valve Sticking

Shortly after resolution of this anomaly, a valve checkout on SV1 (already successfully completed on the other three spacecraft) showed no thrust from any of the four nozzles. The IEP valves used have been observed on multiple occasions to stick closed, especially after long periods of storage, but can typically be freed by repeated actuation. A total of 600 short pulses were commanded, which led to measurable thrust from all four nozzles, although still below the expected levels by a factor of 8. The recovery of all four nozzles at the same time indicated the issue was in the refill valves, as a simultaneous partial unsticking of four separate valves was judged unlikely.

A further round of short cycles, this time including the two refill valves, showed thrust from all four nozzles on the order of 4 mN, in family with the other three systems. A short maneuver was conducted which confirmed this recovery and delivered 10.3 cm/s.

The most likely root cause of this was both refill valves sticking closed due to long storage and possibly exacerbated by unexpectedly cold temperatures post-deployment. One of the refill valves was freed by the first round of repeated cycling, leading to partial pressurization of the plenum and the low thrust results. The second round of debugging freed the second refill valve, providing full thrust.

This could nominally have been confirmed by pressure telemetry from the plenum, which should have shown an increase as the valves became unstuck. However, the plenum pressure sensor op amp has also failed, and no usable pressure data is available. The relatively limited electronics telemetry available on Hamlet has prevented a root cause from being identified, but PCB component failure due to low temperature is plausible.

Assembling the Swarm

The other three spacecraft proceeded through commissioning relatively smoothly. After test burns were conducted, the vehicles performed a series of Drift Control Maneuvers (DCMs) to arrest their drift relative to one another and stabilize their intersatellite distances at 70-100 km. SV2, SV3, and SV4 were now within crosslink radio range and were able to continue with payload checkout and early experiments.

While SV1 was being actively debugged, weekly stationkeeping maneuvers were performed to correct for differential drag, which was primarily caused by periodic safe modes in high drag attitude.

The disturbance force from SV1's leak had reduced its altitude and by early October 2023 it was 3500 km ahead of the other three spacecraft with slight eccentricity and inclination errors. SV1 troubleshooting was still in progress, and given its propellant loss from the leak, the operations team decided to bring the other three spacecraft to SV1.

Three maneuvers one week apart were performed on SV2-4 to lower their altitude and catch up with SV1 over the course of 5 weeks. Two more maneuvers were performed in November to raise altitude to match SV1 and bring it into the same 75-100 km in-track spacing. Shortly after this, SV1 conducted its first successful out-of-plane maneuver, which slightly reduced its inclination error and confirmed that the propulsion system was functioning.

Swarm Operations

Between November 2023 and March 2024, the swarm remained in a consistent in-train formation with maneuvers conducted weekly for stationkeeping. After its final checkout in January, SV1 began participating in these maneuvers. The largest source of disturbance in the formation were the infrequent spacecraft reboots, which caused safe mode entry and a higher drag attitude, requiring the vehicle to raise its orbit after recovery.

Passive Safety Ellipse Entry

In March 2024, the swarm attempted transition into the passive safety ellipse (PSE) formation. This formation produces relative motion between the four spacecraft that is elliptical in the crosstrack directions by varying orbit eccentricity and inclination[7], illustrated in Figure 7.

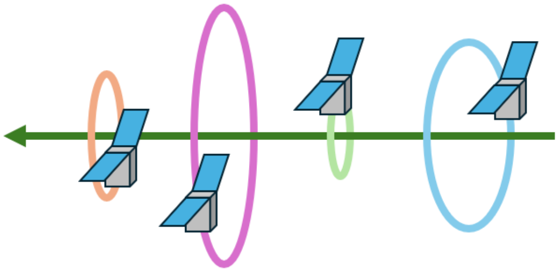

Figure 7: Passive safety ellipse relative motion about reference orbit.

The PSE formation has two advantages. First, the ellipses do not intersect, and so any in-track drift causing two vehicles to pass each other cannot produce a collision. This in-track drift could be caused by poor maneuver performance or an extended period of high-drag safe mode attitude. Second, the PSE produces more apparent relative motion of the spacecraft from the perspective of the rest of the swarm. This relative motion allows the StarFOX experiment to demonstrate more tracking capabilities.

Given persistent comm issues on SV1, it was excluded from the PSE reconfiguration. An unexpected reboot of SV4 shortly before the maneuver caused it to abort, so ultimately only SV2 and SV3 maneuvered.

The reconfiguration consisted of five burns per spacecraft: two prograde, two retrograde, and one crosstrack. Performance variability was significantly worse than normal on these large burns, particularly on SV2, which experienced a range from 41% under target impulse to 7% over target. The net effect of this variability was to cause SV3 and SV2 to switch places in the formation, although since the PSE had been roughly established there was no risk of collision. The drift was arrested by another pair of maneuvers on SV3 and SV2, and the project decided against maneuvering to switch the spacecraft back to the nominal order.

The relative motion achieved was not exactly the intended PSE formation, but produced enough apparent relative motion for StarFOX that further PSE maneuvers were not performed.

Extended Mission

In late May 2024, Starling transitioned from the primary mission to the extended mission, called Starling 1.5. The extended mission focuses on conducting autonomously planned maneuvers in cooperation with SpaceX Starlink satellites for collision avoidance. While these maneuvers are planned onboard, they are downlinked for screening and must be armed via ground command before they can be executed.

As of early June 2024, transition to the extended mission is still underway. During this transition period, SV3 experienced an aborted maneuver due to a nozzle valve sticking closed. Troubleshooting efforts to free this valve are still in progress. Due to the inward angling of the nozzles, 3-DOF attitude control is only possible with all four nozzles functioning. If the valve is not able to be recovered, maneuvering capability will be extremely limited since momentum will quickly accumulate.

PERFORMANCE

Table 1 below shows a summary of the maneuvers performed to date (as of June 2024) on the four flight units. Note that SV1 has accumulated far fewer maneuvers than the other three, due both to its longer commissioning and the spacecraft's ongoing radio issues. SV4 is the second lowest since it did not participate in PSE reconfiguration.

Table 1: Performance summary of the four Hamlet systems as of June 2024

	SV1	SV2	SV3	SV4
Burns performed	16	52	45	37
Total valve cycles	10,674	64,986	47,616	35,988
Delta-V	0.96 m/s	5.44 m/s	4.12 m/s	3.59 m/s
Propellant consumed	29 g 572* g	170 g	128 g	111 g

* SV1 propellant consumption is 572 grams including the leak losses

Thrust Variability

The Hamlet flight units have experienced higher than expected variability in thrust throughout the mission, most notably during the PSE reconfiguration. Hamlet was intended to control thrust by maintaining a consistent system temperature, but thermal variability between maneuvers and slosh dynamics in the main tank have interfered with this.

Despite the similarities between the four spacecraft, there are slight differences in their thermal conditions that lead to different initial conditions for the main tank. Hamlet's thermal conductivity to the bus structure is higher than anticipated, and the main tank heaters are unable to effectively preheat the liquid propellant. This means that the initial conditions are dominated by the bus structure temperature. These conditions can be roughly predicted in advance based on attitude and eclipse timing, but there is still significant unpredictability. There are no temperature sensors at the Hamlet-bus interface, since available temperature sensors were concentrated on valves, sensors, and electronics.

The second source of error is slosh of the liquid propellant in the main tank. Because of the placement of the refill valve inlets and the dynamics of liquid in microgravity, the refill valves will mostly ingest liquid from the main tank, with this liquid boiling to vapor in the manifold, piping, and the plenum itself. This allows the two refill valves to keep up with the mass flow rate required by the four nozzle valves. The temperature of this liquid has a large impact on plenum conditions, since a warm liquid will vaporize more readily and lead to higher temperature and pressure in the plenum.

R-236fa has relatively poor thermal conductivity for a liquid, with a range of 0.074 – 0.082 W/m-K in the typical operating temperature range (c.f. liquid water, with a range of 0.556 – 0.598 W/m-K at similar temperatures). This combined with the lack of convection in microgravity allows significant thermal gradients to build up in the main tank. The constant (albeit small) acceleration during a burn disrupts the liquid, which (clinging to the walls and other surfaces initially) tends to drift towards the "back" of the spacecraft. This causes the warm and cold liquid to mix unpredictably, and the temperature variability at the inlet to the plenum causes large swings in plenum conditions.

Figure 8 below shows a telemetry profile of a ~4 minute burn, showing fairly typical behavior. The gray background indicates the burn time, with plenum pressure in blue. The pressure rapidly drops as the preheated plenum gas is expelled. Refilling from the cold main tank liquid causes the pressure to stabilize at a lower level, before slowly recovering as the main tank liquid mixes and the temperature at the refill valve inlet increases.

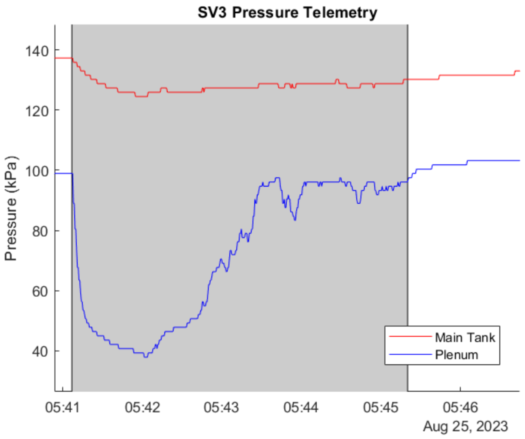

Figure 8: Hamlet telemetry from an early maneuver on SV3, with the burn time backgrounded in gray. Note the plenum pressure (blue) drops suddenly before slowly recovering.

The example given in the figure is broadly typical of the shape of the plenum pressure curve, but the exact curve varies greatly from one maneuver to the next, or even within burns of the same maneuver. In particular, the amount of time required for the pressure to recover varies from 30 seconds to over 5 minutes. The high and low plateau pressures also vary depending on the initial main tank temperature. Attempts at longer preheats involving more heaters were not successful in changing this. A solar preheat, in which the prop system face is pointed towards the Sun, was ruled out since it would leave the spacecraft power negative for too long.

This variation has been coarsely observed in the maneuver reconstruction process. The process is not granular enough to follow these pressure curves in detail, but broadly the beginning of a burn tends to have lower reconstructed acceleration than the end. Burns that are too short for the pressure to recover have more constant values.

The per-nozzle average thrust, as reconstructed from GPS data, tracks very well with the average plenum pressure across the entire burn, as shown in Figure 9. Burns shorter than 30 seconds for which high-rate telemetry was not downlinked are not included, since too few plenum pressure points are available for a reasonable average.

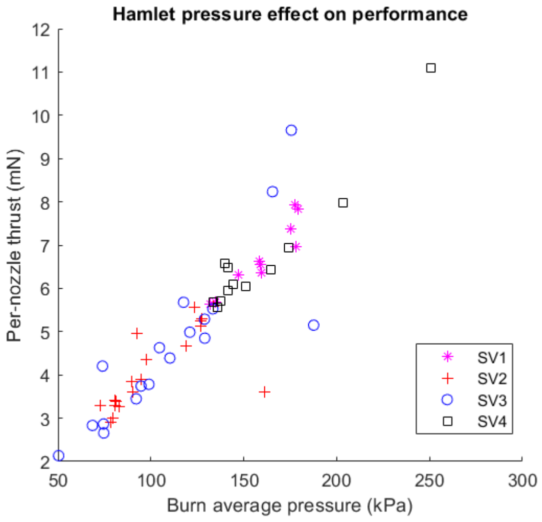

Figure 9: Per-nozzle thrust reconstructed from GPS data compared to plenum average pressure.

Plenum pressure is a very strong predictor of burn-average thrust, but some unexplained variability remains. Ignoring the two low outliers (early aborted burns on SV2 and SV3), the 1-σ deviation from the best fit linear model is ±8.6%. This represents the lowest variability that could be achieved if Hamlet was using a pressure feedback controller, in which input commands were scaled based on current plenum pressure. The current method, which relies on scaling the burn durations on the ground based on predicted conditions, has so far had a 1-σ of ±21.9%, a mild improvement over the deviation against a single nominal thrust value for each unit, which gives a 1-σ of ±30.2%.

Implementing a pressure feedback controller was discussed during development but rejected for two reasons. First, it likely would have required at least modest changes to the BCT legacy thrust controller, incurring development costs. Second, pressure sensor reliability and rapid bias drift was a significant issue on BioSentinel's propulsion system. Hamlet uses different pressure sensors, but there was not sufficient confidence in their flight performance to rely on them for control.

Ultimately, the sensors used on Hamlet have proven extremely reliable. Bias drift (estimated using propellant temperature and saturation pressure) has been under 10 kPa maximum as of June 2024. However, pressure feedback control would have failed on SV1 due to the failed op-amp, requiring some kind of backup controller. Despite this, overall swarm performance would have benefited greatly from pressure-based control, as the fallback solution for SV1 would have been similar to the scheme used currently.

CONCLUSION

Much work remains to be done to streamline Hamlet operations, particularly improving propellant state prediction tools to reduce performance variability. Many lessons have been learned that can hopefully be applied to future fine-maneuvering systems for even more compact swarm formations.

However, Hamlet has largely achieved its role in the Starling mission, performing over 100 total burns and maintaining the swarm formation for the ten-month primary mission. As Starling moves into the Starling 1.5 extended mission, Hamlet will be used as part of the demonstration of autonomous onboard maneuver planning and execution.

Acknowledgments

NASA's Small Spacecraft Technology program have supported and funded the Starling project, including the development of the Hamlet propulsion system. Thanks also go to the Starling integration and test team for helping with critical propulsion system tests. Last but not least, thanks go to the Blue Canyon mission operations personnel for their troubleshooting work on SV1 during spacecraft commissioning.

References

1. Sanchez, H., et. al., "Starling1: Swarm Technology Demonstration," 32rd Annual AIAA/USU Conference on Small Satellites, Logan, UT, 2018.
2. Kruger, J., A.W. Koenig, and S. D'Amico, "Starling Formation-Flying Optical Experiment (StarFOX): System Design and Preflight Verification", Journal of Spacecraft and Rockets, Vol. 60, No. 6. November 2023.
3. Adams, C. et. al., "An Overview of Distributed Spacecraft Autonomy at NASA Ames," 37th Annual Small Satellite Conference, Logan, UT, 2023.
4. Fusco, J., T. Stevenson, and R. Nakamura, "Flight Experience and Lessons Learned From Biosentinel: A 6u Deep Space Cubesat," 46th Annual AAS Guidance, Navigation and Control Conference, 2024.
5. NIST Webbook, "Saturation Properties for 1,1,1,3,3,3-Hexafluoropropane (R236fa)," URL: https://webbook.nist.gov/cgi/cbook.cgi?Name=R-236fa&Units=SI . Accessed May 2024.
6. DSM Functional Materials Group, "Somos PerFORM," URL: www.dsm.com/somos Accessed December 2019.

7. D’Amico, S. and Montenbruck, O., “Proximity Operations of Formation-Flying Spacecraft Using an Eccentricity/Inclination Vector Separation,” Journal of Guidance, Control, and Dynamics, Vol. 29, No. 3, May 2006.