(Preprint) AAS-23-141

ON-BOARD SWARM CONTROL FOR AUTONOMY AND RESPONSIVENESS (OSCAR)

Kenneth Center,* Hanspeter Schaub,† Evan Sneath,‡ and João Vaz Carneiro§

Orbit Logic collaborated with the University of Colorado Boulder (CU) to develop the On-board Swarm Control for Autonomy and Responsiveness (OSCAR) solution. OSCAR leverages Orbit Logic's heritage Autonomous Planning System (APS) onboard planning/response framework and CU's heritage satellite formation flying and orbit control algorithms. Together, these technologies constitute a capability that allows a swarm of planetary-orbiting satellites to dynamically adapt their configuration to accommodate varying mission needs. OSCAR is capable of determining, planning and orchestrating the relative movement of each swarm element to achieve formations that satisfy a variety of needs, including "convoys" allowing events detected by leading satellites to trigger follow-up responses by following satellites, "double-echelons" to allow broad coverage on equatorial passes, or single/multiple synthetic apertures enabling coordinated collection of space-resident or planetary surface data by multiple asset elements.

INTRODUCTION

Addressing some of the needs of future space exploration missions will benefit from robust and effective autonomous collaborative operations between heterogeneous spacecraft teams or swarms. One key challenge in operating these swarms autonomously in environments distant from Earth is realizing reliable automated, cooperative trajectory control that achieves high precision coordinated navigation and control with efficient information exchange, efficient use of onboard resources, and minimal ground commanding. High precision relative positioning and timing is necessary to compose and maintain formations such as sparse apertures that use multiple coordinated spacecraft with small tele-scopes in lieu of a single satellite with a giant telescope. In contrast to approaches which combine data from a single asset across time, including potentially multiple satellite orbits, sparse/synthetic aperture approaches combine data across several satellites, enabling the imaging of rapidly evolving scenes. Imaging applications enabled by this technology include high spatial resolution imaging, 3D thermal imaging, and domains with rapidly evolving phenomena, such as for atmospheric imaging of gas giants. Furthermore, the use of multiple spacecraft provides mission-level resilience in cases where individual nodes are failed or degraded, since objectives can still be satisfied by the remaining healthy elements.

---

* Director of Advanced Programs, Orbit Logic, 1515 Walnut St., Boulder Colorado 80302.
† Professor, Schaden Leadership Chair, Aerospace Eng. Sciences Dpt., Boulder Colorado. Fellow of AIAA and AAS
‡‡ Principal Software Engineer, Orbit Logic, 1515 Walnut St., Boulder Colorado 80302.
§§ Graduate Research Assistant, Aerospace Eng. Sciences Dpt., University of Colorado Boulder.

Orbit Logic and CU worked together to develop a design and initial prototype of the On-board Swarm Control for Autonomy and Responsiveness (OSCAR) solution – with the intention of creating an onboard software solution that allow clusters of satellites to engage in collaborative planning and execution strategies in support of missions where a team of satellites may offer performance, responsiveness, and resilience advantages over a single exquisite spacecraft platform.

The resulting OSCAR prototype has been validated through simulation runs performed against the Basilisk Astrodynamics Framework, developed jointly by the University of Colorado (CU) Autonomous Vehicle Systems (AVS) Lab and the Laboratory for Atmospheric and Space Physics (LASP). Basilisk hosts models of satellite subsystem-oriented functionality representative of previously flown science and exploration missions. The algorithmic logic has also been run on computing elements representative of contemporary satellite flight processors to confirm that the software solution is suitable for execution in constrained processing and memory environments.

SIGNIFICANCE OF THE OSCAR SOLUTION

NASA's future space exploration mission designs suggest multiple concepts employing homogeneous or heterogeneous teams of satellites, including the following two cases:

1. Convoys of spacecraft in which the lead spacecraft triggers detailed measurement of a very dynamic event by the following spacecraft
2. Missions involving distributed space telescopes and distributed synthetic apertures that rely heavily on coordination and control technologies

The "flying formations" of the two use cases above need not be exclusive to a specific mission deployment. Instead, it should be possible for a low-cardinality swarm (on the order of 4 to 15 satellites) to support multiple formation flying configurations according to the needs of the mission at different times or in different phases.

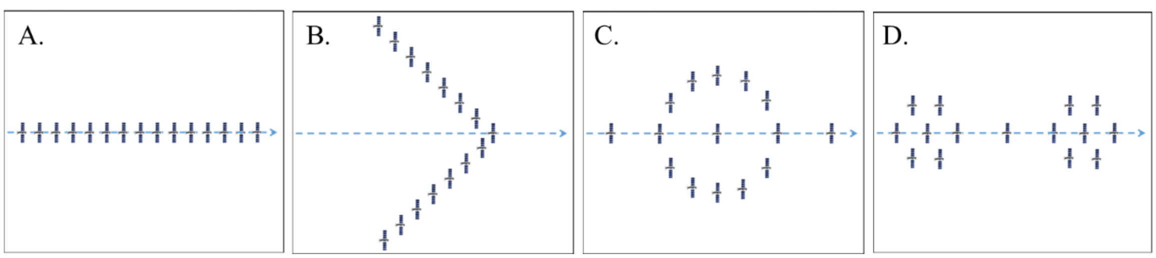

Figure 1. Several prototypical satellite orbital flight formation configurations, such as A. leader-follower, B. staggered swept echelon, C. and D. circular-projected and leader-follower combinations.

Furthermore, it should be possible for the swarm to plan and orchestrate the movement of individual satellites between these configurations without the assistance of mission control. OSCAR's onboard autonomy would determine when one of these formations would be appropriate to meet a specific mission need (in preparation for a specific type of science gathering, for example – or in response to an unanticipated dynamic event), then determine the assignment of individual assets to occupy different "slots" in the next formation and the means by which each asset navigates into that slot, honoring platform level resource constraints (performance limitations, safety keep-outs, preserving consumables) and swarm-level constraints (avoiding collision with other assets during formation transitions). Figure 1 graphically depicts several satellite orbital flight formations relevant to planetary science gathering, where the satellites are moving left to right.

Panel A illustrates a simple "convoy" or "string of pearls" formation. In this formation, satellites toward the front could be performing coarse imaging, sharing the work of fulfilling individual collections of a larger region, or making use of line-of-sight instrumentation such as atmospheric layer profiling to identify "trigger conditions" indicating the likelihood of high-value science. The presentation of these events would result in the tasking of satellites further back in the convoy to employ different instruments (e.g. high-resolution imagers) to capture the episode in greater detail. This could include "time series" collection to record the temporal evolution of the event if relevant. In many cases (depending on the timescale of the event), such a scenario could not be undertaken if mission control was in-the-loop. Coordinated autonomous reaction by the swarm is the only way to ensure timely response to these types of ad-hoc events. Such a formation is relevant to "train of satellite" mission such as A-Train, which uses seven satellites with different sensor instrumentation suites to characterize mid-level clouds on Earth (as described by Kidder, Kankiewicz and Vonder Haar)[1].

Panel B shows another formation – one in which the satellites arrange into an offset time-varying double echelon. This allows a wide orbital swath to be covered twice per orbit (on opposite sides of the central body). Such a formation maximizes simultaneous surface coverage in the regions where the formation is most fanned out, ensuring that data collected during the passes are fairly closely time synchronized. In cases where this is important (e.g. characterizing weather at the planetary day/night terminator) this type of formation is most useful. Twice an orbit, during crossing of the satellites' orbital paths, the staggering of the echelons allows satellites to "thread-through" without collision risk.

Panel C illustrates a synthetic aperture formation using a combination of circularly projected relative orbits combined with leader-followers. Organizing into such a formation allows for certain assets to serve as discrete elements of a very large antenna, with one asset dedicated to act as the "feed" node that receives and combines the collected signals into an integrated data product. Alternatively, Bekey[2] describes particular cartwheel relative orbits that allow the circular relative radius pattern to be preserved (albeit with a slow rotation of the "antenna" perimeter) over the entire orbit with minimal thrust corrections, capitalizing in natural orbital motion with small differences in each orbit's inclination and eccentricity. This formation additionally includes a single satellite in the lead position (ahead of the synthetic aperture) with might be used to detect signatures of interest to cue the acquisition of data using the aperture. A "spare" asset (at the rear of the formation) is additionally available to be "substituted in" to the aperture formation should one of the active elements degrade.

Finally, panel D shows a variant in which two synthetic apertures are formed. This would allow two collection activities to be performed separately and in parallel. This formation would be useful in situations where the resolution offered by a synthetic aperture composed of the majority or the available satellite nodes is not required, but there is a desire to increase the tempo of collections. As in the single synthetic aperture case, a single node is positioned between the two apertures. It could be used as a communication relay (if the two apertures are spaced significantly far apart that they cannot communicate directly), as a detection and trigger node to tip collection by one of the two apertures, or as a hot spare (as described in the previous paragraph).

The primary focus of the OSCAR activity was to develop and validate an initial software prototype that would allow a cooperating cluster of satellites to dynamically plan and execute transitions between flying formations in response to dynamically detected events.

SUPPORTING TECHNOLOGIES

Our team brought together two key technologies to create the initial OSCAR design and prototype. Firstly, we leveraged Orbit Logic's Autonomous Planning System (APS) to serve as the decentralized planning framework for OSCAR. Secondly, we leveraged the University of Colorado's Basilisk framework* for the spacecraft's functional subsystem interfaces and to support dynamic satellite simulation in planetary environments.

Orbit Logic's Autonomous Planning System (APS)

APS (see Figure 2) is on-board autonomy software for satellites or other unmanned robotic assets[3,4,5,6]. It enables the planning and execution of missions onboard the satellite, taking only high-level directives from ground-based operators. Onboard execution allows satellites to react more quickly to on-orbit identified events, because the delays associated with communications to the ground and operator decision making are completely removed. Moreover, this makes systems more robust to scenarios where communications are interrupted, or operators are overtasked; APS also addresses the fact that the ground station can potentially be a single-point-of-failure. APS employs modular Specialized Autonomous Planning Agents to construct mission activities to satisfy high-level mission goals/objectives. The Master Autonomous Planning Agent maps the activities of multiple SAPAs into a figure-of-merit optimized, deconflicted plan for the use of local platform resources. This plan can be passed along to the hosting system's native command timeliner, or (optionally) to APS's Mission Timeliner module. The APS Vehicle Interface Translator is a highly configurable component that maps between the hosting system's native data protocols and the APS architecture's message formats. It allows APS to receive data from the hosting platform's C&DH subsystem (time, vehicle state, subsystem status, sensor data, and other local telemetry sources), user requests or configuration settings from the ground system, share teaming data with other APS-enabled assets (through the hosting platform's communication resources), and deliver resource commands to the hosting satellite to realize autonomous operations.

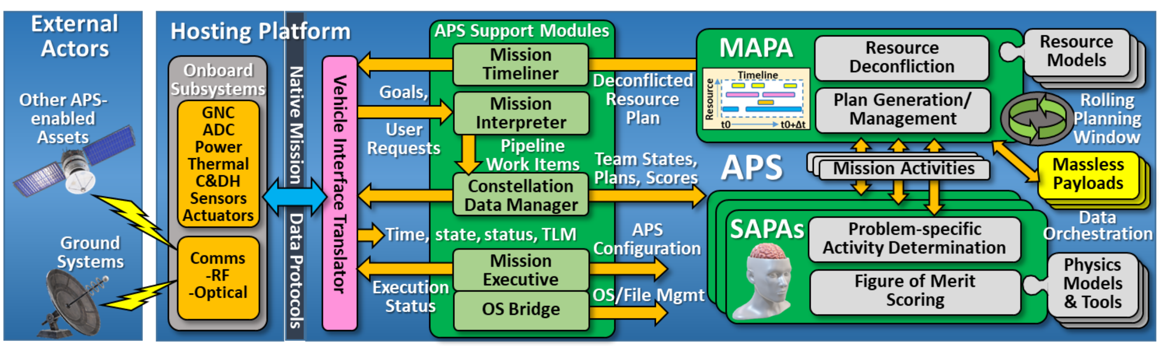

Figure 2. APS for Asset- and Team-level Autonomy.

The APS architecture is constructed such that it is possible to utilize inter-asset communication links to share data in support of constellation- or swarm-level autonomy – even when those links have low-availability or capacity. The Constellation Data Manager module facilitates this information sharing by maintaining "work pipelines" describing the relationships between "work items" representing steps in the TCPED process (Tasking, Collection, Processing, Exploitation, and Dissemination). Instances of APS on each participating asset consider their ability/opportunity to perform these work items, using a decentralized bid/pass strategy to determine which asset will execute its plan (based on the highest multi-factor figure of merit). An onboard Mission Interpreter

* https://hanspeterschaub.info/basilisk

module generates work pipeline representations from high-level user data requests. APS's fully decentralized autonomous planning capability enables cooperation and teaming between assets to dynamically replan in response to events detected by the satellites themselves, which are frequently unanticipated. If inter-satellite links are down or unavailable, APS replans independently for each asset, adjusting task execution schedules to react to ad-hoc on-orbit events, such as the emergence of a fault. If inter-satellite communication links are available, APS will use them to coordinate responses across the constellation team – for example, to shift the tasking load among satellites to ensure coverage and limit redundancy. APS allows third party developers to integrate custom algorithms (or "Massless Payloads") and specify how they are invoked as part of the TCPED process for given mission workflows. The APS Mission Executive manages a specified collection of modular applications during a mission phase – assessing the health of each application and restarting those detected as failed or unresponsive. The plug-and-play middleware allows restarted applications to re-join the data system immediately and without service interruption.

APS is currently flying in space on the YAM-3 satellite as part of a Space Development Agency (SDA) demonstration mission for onboard plan optimization and data processing orchestration in support of a maritime domain awareness mission. In addition, APS has been applied to satellite self-protection and space-based local Space Situational Awareness (SSA) in the AFRL SaFIRE program, distributed collection planning in the DARPA Blackjack Pitboss program, and collaborative autonomous patrol/detect/track with unmanned underwater vehicles swarms in a Navy program.

Basilisk Astrodynamics Simulation Framework

The spacecraft relative orbit control and targeting pointing simulation work for the OSCAR effort was based on the novel open-source Basilisk astrodynamics simulation framework[*]. As the example shows in Figure 3, it provides a highly modular simulation environment using a message passing interface to connect simulation modules and components. Basilisk is being jointly developed with LASP to support the ADCS sub-system of the recently launched Emirate Mars Mission Hope[†] and had been used previous to that on a range of small satellite missions. It is available as an open-source framework and used by research labs and industry world-wide. The physics modeling is flight-validated, and the modular flight software components are employed in missions. Significant work was performed in the course of the OSCAR effort to facilitate the creation of scalable multi-satellite simulation scenarios - with rapid prototyping and modularity as core principles[‡].

Vizard is a companion software to Basilisk which allows for the Basilisk simulation to be visualized in a three-dimensional environment using the Unity Gaming engine. The simulation setup will be configured to stream the simulation states to Vizard for mission operation validation.

[*] https://hanspeterschaub.info/basilisk
[†] https://emiratesmarsmission.ae
[‡] https://hanspeterschaub.info/Papers/Carneiro2022a.pdf

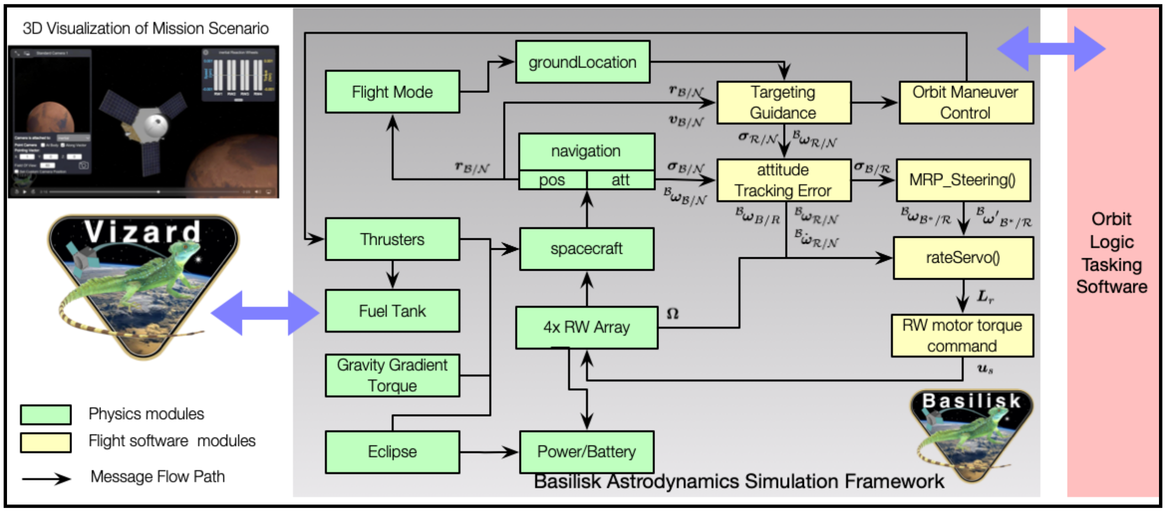

Figure 3. Basilisk Simulation Workflow for Commanded Pointing and Maneuvering.

OSCAR ACHITECTURE

The OSCAR team designed and implemented a modularity-focused design and assembled a modeling and simulation environment capable of validating the key capabilities of the solution. All new work was implemented to modular service architecture standards (we use AFRL's Aspire[6] middleware to interface between APS modules, and protobufs* over gRPC† to interface the mission planning solution to spacecraft-hosted resources).

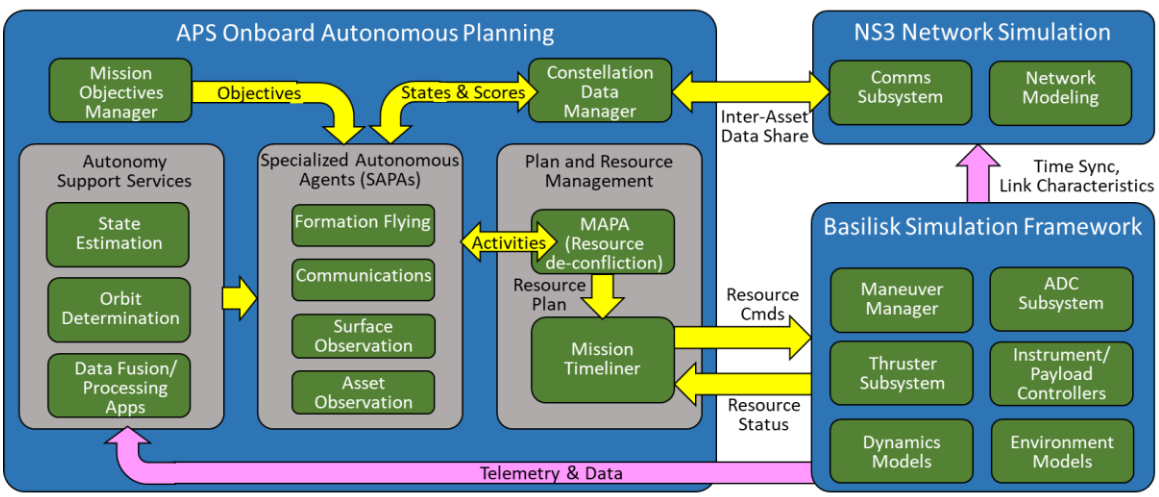

Figure 4. APS extended for the OSCAR solution.

Figure 4 depicts the APS-based architecture that resulted from our OSCAR prototyping effort. It derived heavily from existing and related work. As such, a significant portion of the software code in the diagram already existed and merely required integration and configuration to reach a state where meaningful functional demonstrations could be conducted. The major areas of work were as follows:

---

* https://developers.google.com/protocol-buffers/docs/reference/overview

† https://grpc.io/docs/guides/

1. We extended the Constellation Data Manager module to include the ability to decompose team-level objectives enabled by formation flying configurations into work items that could be used by the APS planning framework to determine asset-level allocations to meet those objectives. This involved additions to the database structure of the CDM to manage those mappings and definitions.
2. We worked with CU to leverage/extend code for the "Formation Flying SAPA" shown in the *Specialized Autonomous Agents* box of the diagram. This allowed the APS planning system to determine satellite formations that are appropriate for meeting certain mission objectives as asserted by the *Objectives Manager*. This new SAPA considers the states/orbits of the individual assets in the team, determines assignments to occupy "slots" or "roles" in the next formation and determines the maneuvers (combination of attitude control and the application of thrusters) required to complete the transition.
3. To plan the maneuvering behaviors required for orbit adjustments of the formation spacecraft, a new "activity class" was developed. This allowed the Formation Flying SAPA module to engage in asset-level resource deconfliction with the MAPA module.
4. We developed a new resource plug-in for the MAPA module to facilitate the commanding and status associated with the use of a propulsion resource. This allowed deconflicted plans generated by the MAPA for satellite orbital maneuvering to be managed by the Mission Timeliner module.
5. Orbit Logic worked with CU to create a custom translator to map between the APS message standard and the Basilisk message protocols. This allowed APS to interoperate with Basilisk in a real-time mode to acquire satellite resource state/status and issue resource-level commands to the simulated satellite subsystems that Basilisk models. Attitude control, thrusters, and candidate instruments/payloads were all connected for the initial prototype development and validation effort. Completion of these extensions facilitated real-time system simulations involving hardware in-the-loop.

Computing, Specifying and Maintaining Formation Orbits

As part of the software prototyping effort, CU performed work to leverage their orbital mechanics expertise and existing Basilisk framework to arrive at a collection of algorithms that could determine and orchestrate the transitions between the three primary formations (string-of-pearls, double echelon, and sparse aperture) that the team identified in the proposal as being relevant to future exploration missions. Getting a capability in place to support the OSCAR effort involved determining how to parameterize each of these formation types (in terms of delta orbital element specifications) and then how to perform a series of deltaV maneuvers to transition from one orbit of the initial formation to another orbit in the next formation.

The relative motion of the orbiting spacecraft will be controlled using a near-fuel optional impulsive feedback control strategy using orbit element differences[7,8]. The satellites are assumed to be coordinating through continuous communication as shown in Figure 5. With the impulsive relative orbit control strategy, the Cartesian inertial navigation solution of the position and velocity vectors $X_1 = (r_1, v_1)$ and $X_2 = (r_2, v_2)$ are mapped into the corresponding osculating orbit elements oe_1 and oe_2. Next the relative motion is evaluated as the orbit element difference:

$$\Delta oe = oe_2 - oe_2 \tag{1}$$

Assume the classical orbit element set $oe = (a, e, i, \Omega, \omega, M)$ is used, then the first 5 elements are invariants of the Keplerian motion. The mean anomaly angle M is the only term that varies with time. Given a set of desired differential orbit elements Δoe^*, the relative orbit trajectory tracking error is then defined as:

$$\delta oe = \Delta oe - \Delta oe^* \tag{2}$$

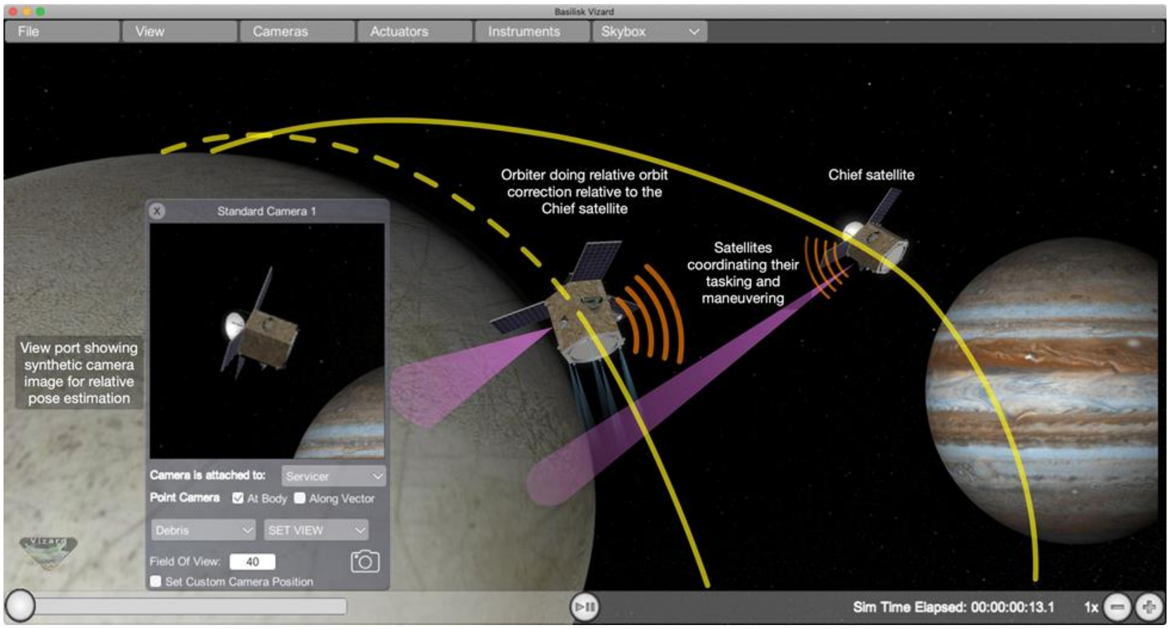

Figure 5. Basilisk applied to OSCAR solution.

The control strategy to be employed is as follows:

1. At a specified orbital epoch time, determine the relative trajectory tracking error δoe
2. Determine a firing sequence to apply over the following orbit including[7]
    a. Burns at apoapsis and periapsis to correct sets of (a, e) and (ω, M) tracking errors
    b. A single burn at the orbit plane intersection to correct the (i, Ω) out of plane motion
3. Repeat this process every orbit for closed loop station keeping or perform this once for a single large relative trajectory maneuver.

While this control law is near fuel-optimal, instantaneous burns are not physically feasible, and the simulation treats each burn as continuous for a period of time. The thrusters are designed to be capable of delivering the necessary delta-V in a short amount of time (minutes), such that the instantaneous burn assumption is viable. Further, the feedback nature of this control allows for effective station keeping. The relative motion simulation includes both translational and rotational motion of each spacecraft. This yields a high-fidelity simulation where to achieve a desired Δv orbit correction, the spacecraft orientation must be rotated and controlled such that the thrusters are pointing in the desired correction during the burn. Post-burn the orientation should return to the target pointing mode.

Constellation design was a significant part of this work, and each formation was designed using relative orbital elements differences. Three formations were explored: string of pearls, double echelon, and circular aperture, each having a specific advantage for data collection.

The string of pearls formation is used to collect data on a target for a longer period. Since all spacecraft are following each other in a line, the target is kept in the line-of-sight of at least on spacecraft for longer. In terms of orbital elements, the relative orbits of each spacecraft with respect to the chief are defined by a difference in true anomaly, while keeping all other orbital elements the same. For this work, a difference of 0.1º between each satellite is used.

The double echelon formation is useful for sweeping a large area. The horizontal spacing between spacecraft allows each instrument to not overlap with each other, guaranteeing a better coverage of the surface of the planet. The relative orbit elements are defined in pairs (one on each side of the echelon). A true anomaly difference of 0.1º and a right ascension of the ascending node difference of ±0.1º are used to define each spacecraft pair along the echelon. However, this is not enough to define the double echelon formation, due to the coupling between the right ascension of the ascending node (RAAN) and the argument of periapsis (AoP). Therefore, depending on the value of the RAAN, a correction to the AoP is added[7].

The circular aperture is used to gather higher resolution pictures of a point on the planet, as the spacecraft are assembled such that they mimic a larger lens. To define the relative formation, the Clohessy-Wiltshire approximation is used. The circular shape is defined in the Hill frame, which is then converted to the inertial frame. Once the cartesian differences in the inertial frame are defined, they are converted to orbital elements, which are used to define the relative orbital elements. The resulting orbital element differences consist of some eccentricity difference but are mostly dominated by changes in true anomaly and argument of periapsis. For this formation to work, the chief's orbit must be almost circular (eccentricity of 10^{-6}). The changes in true anomaly and AoP are quite large, although their sum is always close to 0. Because the chief's orbit is quasi-circular, this results in a small relative distance between each satellite.

Decentralized Planning Approach

An OSCAR-supported scenario starts with formation objectives being supplied to APS. These objectives generally state that a certain formation pattern should be used in certain mission situations (e.g. a double-echelon formation should be used when the multi-spectral imager is being used for wider area search, or that the string-of-pearls formation is to be used when the narrow field of view imager is being used for high-cadence temporal collects). These objectives are decomposed into a collection of formation-relevant "work items" that include relative orbital slots that need to be occupied to satisfy the formation pattern. Once these entries exist, APS employs its SAPA modules to plan optimized assignment of satellites to these slots when different mission trigger events indicate the need to transition to a certain formation flying pattern.

The sequence of data interactions that occurs in the OSCAR architecture for a typical case are:

1. An objective related to a flying formation is received from a mission operator or in response to some mission trigger event. The objective states the need for a specific formation type, that formation's characteristics, and potentially temporal constraints on when the formation must be assembled to support the mission need.
2. The objective is decomposed to determine orbit slots (individual orbits that each spacecraft might occupy to construct a new formation) and create work items representing the need to plan them
3. The viability and cost of transitioning "ownship" from the its current orbital state to each possible slot in the new formation is assessed (multi-factor figure-of-merit, or FoM scores)
4. These scores are posted back to the team via the CDM decentralized database

5. Decentralized logic is employed to determining which satellite will be designated as the "formation leader" (based upon a deterministic heuristic)
6. The Formation SAPA on the leader instance of APS in invoked, which considers all permutations of asset-to-slots in the proposed formation and selects an assignment set that achieves the lowest aggregate FoM-based team score
7. These assignments are posted to the cluster via the CDM
8. Each asset is informed about its slot transition assignment via CDM status changes
9. The resource utilization schedule (attitude control commands and thruster firings) is acquired and incorporated into the APS master timeline
10. The Maneuver Manager is asked to execute the selected maneuver plan (the one corresponding to the assignment determined by the Formation Flying SAPA
11. Execution is orchestrated by the maneuver Manager while in progress
12. If a maneuver plan change is necessary (including cancelations), the deltas are reported back to APS
13. APS reflects the changes in its managed timeline
14. Other activities needing to use the attitude or propulsion resources (or any other limited resources such as power) will re-plan around the updated timeline

DESIGN REFERENCE MISSION DESCRIPTION

The Venus Design Reference Mission (DRM) document from the 2018 Workshop on Autonomy for Future NASA Science Missions[9] describes two DRMs, a simpler nearer-term mission in the 2023-2032 timeframe, and a more complex mission in the 2033-2042 timeframe. DRM1 cites the use of a large orbiter and a number of federated smaller spacecraft whose primary role is to acquire gravity, topography and spectral imaging data to create maps, but also to seek out volcanic/seismic activity. These assets are presumed to be networked together (with each other but also with other classes of assets including atmospheric vehicles and surface probes). Our OSCAR technology would specifically target the coordination between these satellite elements to manage their orbits in ways that benefit the collection of data by the spacecraft, but also enhance the ability of the platforms to exchange data over their comms links to process and fuse it (since some assets may have more capable computing and memory resources than others).

The report goes on to state that "for more complex missions with multiple vehicles, autonomous systems enable the collection and correlation of data from the same phenomena observed from different vantage points to potentially identify instantaneous events—such as erupting volcanoes and Venus-quakes. Monitoring such events over time is needed to discern patterns. Leveraging advances in automation and autonomy can significantly broaden future Venus scientific discoveries." This statement beckons for the capabilities that a mature OSCAR solution could provide. Spacecraft in static orbits with coordinated onboard autonomy could support these objectives, but would do so opportunistically (given relative orbital geometry). OSCAR would allow the spacecraft to work together to position themselves in orbits most favorable to supporting certain mission capabilities. Obviously, the price of this capability is periodic use of propulsion systems to achieve relative maneuvering, but paying that price could potentially result in realizing enhanced science value.

For our OSCAR prototyping effort, we assumed a team of 3-15 small satellites in a coordinated orbit around Venus. Using propulsion subsystems, they would be able to transition between (and maintain) multiple formations in support of mission-enhancing team behaviors.

The two formations are "string of pearls" (all satellites in the same orbit spaced out along the orbit track by mean anomaly or specified separation distance), and "double echelon" (a staggered 'v' shaped formation where the "leader" is the tip of the v and other satellites are behind and off to either side of the leader's orbit. The staggering along the orbit track allows following satellites to thread between each other during node crossing that happen twice an orbit.

The double echelon formation is used for broad area search, so that the team can cover a maximum swath of the planet or moon's surface on each pass with its wide field of view (WFOV) sensors. Processing of the WFOV sensor data may reveal areas of specific interest (in the case of Venus, WFOV IR sensors would be used to seek active volcanic eruptions).

If active eruptions are detected, the desire is to maximize the temporal acquisition of visible band imagery, to capture the atmospheric effects of the eruption using narrow field of view sensors (NFOV).

The "string of pearls" formation is used for sequenced collection of NFOV imagery from a similar orbital vantage point, but spaced over a period of time. As soon as a WFOV detection occurs, the OSCAR solution should plan and orchestrate a transition from the double echelon formation to the string of pearls formation in preparation for temporal collection (hopefully on the next orbital pass of the object of interest).

The subsections to follow discuss aspects of the configuration of this Venus-centric scenario, with emphasis on the satellite formation’s concept of operations.

Satellite Orbits

Table 1 describes the orbits that are being specified for our initial scenarios. Since this is formation flying, we identify a “root” spacecraft (“chief” or “leader”) that serves as the anchor for the formation, whose orbital characteristics can be specified, and then for different formation patterns (string of pearls and echelon) we define delta orbital elements specifications for each of the other formation satellites. Our initial scenarios are constrained to 3 spacecraft (a leader and two followers) to simplify verification of capabilities.

Table 1. OSCAR Validation Orbit Assumptions.

“Root” Mission Orbit	• **Altitude: 100km circular (Semi-major/Semi-minor axis = 7,052km)** • **Inclination: 45deg**
String of Pearls Formation	• Satellites separated along orbit by 5 deg True Anomaly (TA)
Echelon Configuration	• Satellites separated along orbit in increments of 1 deg True Anomaly (TA = -1 for follower 1 and -2 for follower 2) • Satellites occupy adjacent off-track orbits that deviate from the root by 2 deg relative to the Right Ascension of the Ascending Node (RAAN) (+2 deg RAAN for follower 1 and -2 deg RAAN for follower 2)

Spacecraft Configuration

In addition to requesting details on possible missions for OSCAR, we also needed to define a set of characteristics for the spacecraft that might be utilized to support those future missions. In order to host the types of science instruments of the caliber needed to perform the missions in 2018 Venus DRMs, larger busses would be necessary. We performed a system engineering activity to design a 300-400kg spacecraft capable of hosting several exquisite sensor payloads while also offering the thruster nozzles and propellant mass to provide significant orbital maneuvering capability over a mission of reasonable lifetime (1-2 years). Table 2 below cites our initial assumptions about the overall properties and subsystem-oriented capability of a spacecraft that might participate in an OSCAR swarm. These characteristics can easily be adjusted in the OSCAR solution's modeling and simulation environment.

Table 2. OSCAR Validation Assumed Spacecraft Characteristics.

Gross Physical Characteristics		
	Mass Properties	• Dry mass: 250kg • Propellant mass: 100kg • BOL spacecraft wet mass: 350kg • Inertia Matrix (only diagonal elements): (100 ,100 ,55) kg*m^2
Attitude Control		• 3-wheel orthogonal reaction wheel assembly
	Reaction Wheels	• Flywheel Moment of Inertial: 0.191 kg*m^2 • Max Torque Authority: 0.212 N*m • Momentum Storage Capacity: 100 N*m*s
Power Subsystem		• Assume a bus voltage of 28V nominal
	Aggregated Bus Loads	• Idle: 150W (assumes regular communication with other team members over RF links) • Power Addition to During Active Attitude Slews: 200W • Power Addition while Imaging/Processing: 50W • Max load: 400W
	Solar Arrays	• Cell efficiency: 30% • Array area: 4 m^2 • Note: 2 or 4 panel wings, all fixed in body frame and oriented in same direction • Note: Solar Irradiance in Venus orbit is 2200 W/m^2 • Max Collection Power: 2,640 W
	Battery	• Capacity 600 W*hr
Propulsion Subsystem		• Lifetime: 100kg of fuel would be able to supply thrust for duration of 24 hours over the course of the mission (based on calculations).
	Thruster	• Max thrust: 25N • Isp: 215 s • Based on Emirates Mars mission characteristics
	Propellant Tank	• Capacity to hold 100kg of propellant
Mission Sensors		
	WFOV IR Sensor	• Frame camera fixed to the satellite body frame • FOV of 52x52 deg (full angle) provides coverage of an area 1000x1000km on the planet's surface
	NFOV Optical Imaging Sensor	• Frame camera fixed to the satellite body frame • FOV of 2.3x2.3 deg (full angle) provides coverage of an area 20x20km on the planet's surface
Bus Packaging		• Solar Panels are oriented to -Z in the satellite body frame • The boresights of both imagers (WFOV and NFOV) are co-oriented to +Z in the satellite body frame • The Thruster is oriented to -Z in the satellite body frame

Supplemental Sparse Aperture Formation

While the Venus DRM only utilizes the string-of-pearls and double-echelon flying formations, our team was interested in validating the ability of the OSCAR solution to handle the more challenging case of transitioning to and from a sparse aperture formation, which is a very strictly defined collection of orbits that result in N-1 spacecraft forming a ring around a central spacecraft in a root orbit (acting as the feed element that receives and processes all collected signals). This ring is circular when projected on a planet's surface and the spacecraft that comprise the ring circumnavigate the central feed node once per orbit. The results in the next section therefore represent a series of tests in which transitions are made between all three formation types.

RESULTS

With the completed prototype, the OSCAR team created scenarios to show that the following sequence could be fully realized by the decentralized APS onboard planning:

1. Begin with the satellite cluster in the sparse aperture formation (with the presumption that a mission activity is initially in progress where the satellites are performing synchronized collections of space-resident or surface-resident targets, then delivering those data products to the central ("root" or "chief") satellite for processing that results in a second-stage, higher-fidelity product.
2. Transition to string-of-pearls formation. Something in the processing of the sparse aperture data (or some other trigger event) has identified the need for the team to assume a new formation that allows for "tip and cue" behaviors where the satellites at the lead of the formation can perform broad area survey and "tip" the satellites toward the back to "pick up" the event with more targeted sensing approaches.
3. Transition from string-of-pearls to double-echelon. Yet again, a trigger event determined in the course of collection and processing has identified the need to change the team's formation to be more advantageously configured for wide coverage on an orbit pass.
4. Throughout steps 1-3, collections on surface targets not related to the dynamic formation adaptations (nominal collections that are part of the "mission backlog") are being planned and executed while in settled formations, but also during the transitions.

In the screenshot mosaic of Figure 6 (sequence clockwise from upper-left), the initial formation of seven satellites is in the sparse aperture formation, which has the characteristic of six of the satellites circumnavigating about the "root" satellite once per orbit in an elliptical ring that from the surface of the planet appears perfectly circular. Transitioning the satellites forming the ring to the new string-of-pearls formation involves several delta-V maneuvers, which can be seen to occur in the evolving sequence of the figure. Finally, all satellites are arranged in a line following the "chief", whose orbit has not changed, but whose role had changed from "sparse aperture feed node" to "string of pearls leader".

In the screenshot mosaic of Figure 7 (sequence clockwise from upper-left as before), the string-of-pearls formation transitions to the double-echelon formation. Because there is a staggering of the echelon to facilitate collision-free transits of each wing twice an orbit, one side of the formation must first shit along the orbit "drag" direction and then enter a second maneuver sequence to deploy out to each satellite's swept positions.

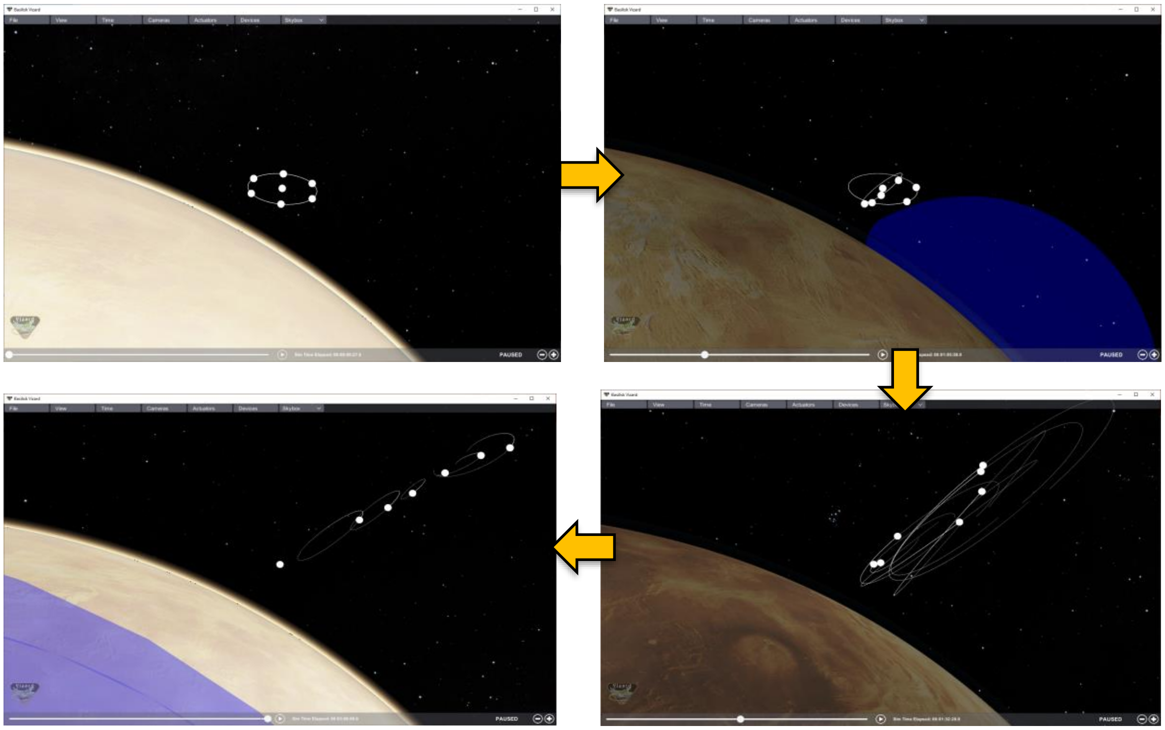

Figure 6. Screenshot sequence of transition from sparse aperture to string-of-pearls formation.

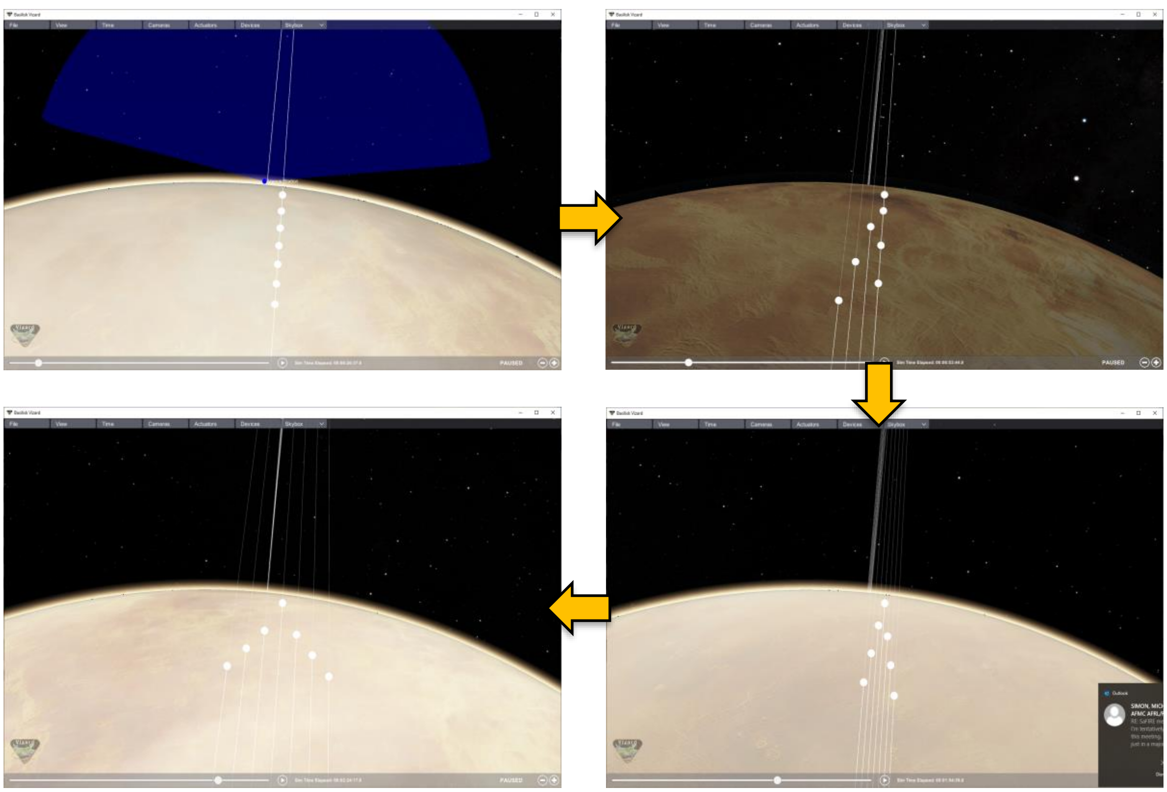

Figure 7. Screenshot sequence of transition from string-of-pearls to double-echelon formation.

As previously mentioned, throughout these orbit transitions the APS planning was determining when it was possible for individual satellites to perform surface target collections associated with standing “user orders” and scheduling those actions at times when they would not conflict with the use of the attitude adjustments and thruster firing sequences facilitating the delta-V maneuvers being used to transition between the various formations.

The thrust of the described test scenario was to verify that we had created an OSCAR prototype system capable of demonstrating a range of mission-relevant, event-driven, adaptive cluster formation use cases – which we did for tests employing up to seven satellites. The exercise proved that transitions between multiple formations of unique mission value can be triggered by conditions sensed by any of the federated satellites, planned very rapidly, and then completed (in most cases) within a single orbit’s time. This capability could enable a wealth of autonomous conops for future missions that might significantly increase science value and potentially even enable the capture of science data that would otherwise be missed.

CONCLUSION

The Orbit Logic / CU Team successfully defined an approach and plan to bring together our respective technologies (the APS planning framework on the Orbit Logic side, and the Basilisk framework on the CU side) to form an integrated and effective solution for dynamic, adaptive autonomous formation flying control of satellite clusters. CU enhanced Basilisk’s ability to configure and model satellite maneuvering in a chosen planetary or moon environment, and was instrumental in deriving formal specification approaches for the various formations (string-of-pearls, double-echelon, and sparse aperture) that the team selected as test cases for the proof-of-concept OSCAR prototype. Orbit Logic successfully extended APS to manage the formation specifications and facilitate the decentralized planning and orchestration required to dynamically trigger transitions to new formation, ensuring that individual satellite resource constraints are honored. These capabilities were validated in a series of simulation scenarios executed against CU’s Basilisk framework. Orbit Logic additionally verified that the onboard software code was executable withing the processing and memory footprint of computing elements representative of contemporary spacecraft mission flight processors.

ACKNOWLEDGMENTS

The IN-PASS team would like to acknowledge NASA for its support of this research under STTR Contract 80NSSC21C0117, and in particular thank our technical monitor Dr. Joshua Vander Hook at the Jet Propulsion Laboratory for his oversight and insights.

REFERENCES

[1] S. Kidder, J. Kankiewicz and T. Vonder Haar. "The A-Train: How Formation Flying is Transforming Remote Sensing," https://atrain.nasa.gov/publications.php.

[2] I. Bekey, “Formation Flying Picosat Swarms for Forming Extremely Large Apertures,” Proceedings of the ESA/ESTEC Workshop on Innovative System Concepts, Noordwijk, The Netherlands, 21 February 2006.

[3] N. Dhingra, K. Center, E. Herz, E. Sneath, S. Gagnard, “APS: Multi-Domain Decentralized Planning for Responsive Multi-Asset Collaborative Autonomy,” Proceedings of the 16th International Conference on Space Operations, Cape Town, South Africa, 3-5 May 2021.

[4] K. Center, E. Herz, S. Gagnard, R. Glissmann, R, Juarez, N. Dhingra, “Autonomous Planning System (APS) for an Onboard TCPED Pipeline,”. Proceedings of the 2021 Small Satellite Conference, Logan, UT, August 9-12, 2021.

[5] E. Herz, J. Neel, M. Simpson, D. George, T. Esposito, K. Center, "Onboard Autonomous Planning System," AIAA Space Ops Conference, May, 2014.

[6] K. Center, P. Courtney, R. Adams, D. Musliner, M. Pelican, J. Hamell, D. Kortenkamp, M.B. Hudson, J. Fausz, P. Zetocha, "Improving Decision Support Systems Through Development of a Modular Autonomy Architecture," Proceedings of the 2012 I_SAIRIS Conference, Turin, Italy, Sept 2012.

[7] H. Schaub and K. T. Alfriend, "Impulsive Feedback Control to Establish Specific Mean Orbit Elements of Spacecraft Formations", AIAA Journal of Guidance, Control and Dynamics, Vol. 24, No. 4, July–Aug., 2001, pp. 739–745. doi:10.2514/2.4774

[8] H. Schaub and J. L. Junkins, Analytical Mechanics of Space Systems, AIAA Education Series, Reston, VA, 2003. doi:10.2514/4.861550

[9] 2018 Workshop on Autonomy for Future NASA Science Missions: Venus Design Reference Mission Reports. https://science.nasa.gov/technology/2018-autonomy-workshop/