NASA/TP-20260005089

Multi-Agent Swarm State of the Art Report

Caleb A Adams
Trey Smith
Aaron Woodard
Brian C Kempa
Jeremy D Frank
NASA Ames Research Center, Intelligent Systems Division, Mountain View, California

Ellemieke Van Kints
KBR, Houston, Texas

June 2026

NASA STI Program Report Series

Since its founding, NASA has been dedicated to the advancement of aeronautics and space science. The NASA scientific and technical information (STI) program plays a key part in helping NASA maintain this important role.

The NASA STI program operates under the auspices of the Agency Chief Information Officer. It collects, organizes, provides for archiving, and disseminates NASA's STI. The NASA STI program provides access to the NTRS Registered and its public interface, the NASA Technical Reports Server, thus providing one of the largest collections of aeronautical and space science STI in the world. Results are published in both non-NASA channels and by NASA in the NASA STI Report Series, which includes the following report types:

- TECHNICAL PUBLICATION. Reports of completed research or a major significant phase of research that present the results of NASA Programs and include extensive data or theoretical analysis. Includes compilations of significant scientific and technical data and information deemed to be of continuing reference value. NASA counterpart of peer-reviewed formal professional papers but has less stringent limitations on manuscript length and extent of graphic presentations.

- TECHNICAL MEMORANDUM. Scientific and technical findings that are preliminary or of specialized interest, e.g., quick release reports, working papers, and bibliographies that contain minimal annotation. Does not contain extensive analysis.

- CONTRACTOR REPORT. Scientific and technical findings by NASA-sponsored contractors and grantees.

- CONFERENCE PUBLICATION. Collected papers from scientific and technical conferences, symposia, seminars, or other meetings sponsored or co-sponsored by NASA.

- SPECIAL PUBLICATION. Scientific, technical, or historical information from NASA programs, projects, and missions, often concerned with subjects having substantial public interest.

- TECHNICAL TRANSLATION. English-language translations of foreign scientific and technical material pertinent to NASA's mission.

Specialized services also include organizing and publishing research results, distributing specialized research announcements and feeds, providing information desk and personal search support, and enabling data exchange services.

For more information about the NASA STI program, see the following:

- Access the NASA STI program home page at http://www.sti.nasa.gov

NASA/TP-20260005089

Multi-Agent Swarm State of the Art Report

Caleb A Adams
Trey Smith
Aaron Woodard
Brian C Kempa
Jeremy D Frank
NASA Ames Research Center, Intelligent Systems Division, Mountain View, California

Ellemieke Van Kints
KBR, Houston, Texas

National Aeronautics and
Space Administration

Ames Research Center
Moffett Field, CA 94035-1000

June 2026

Available from:

NASA STI Program / Mail Stop 050
NASA Langley Research Center
Hampton, VA 23681-2199

Table of Contents

Introduction

The Next-Generation Multi-Agent Swarm (NGS) Study conducted by NASA's Ames Research Center for NASA's Space Technology Mission Directorate (STMD) will develop a comprehensive understanding of emerging multi-agent swarm capabilities. The study aims to identify existing swarm capabilities and asses their potential for persistent lunar space situational awareness, surface monitoring, and distributed autonomy demonstrations. A key objective is to inform the design of a next-generation multi-agent swarm that can perform autonomous distributed remote sensing, position, navigation, and timing (PNT) services, automated deployment that leverages autonomy, edge computing, and interoperable networking to enable cooperative operations without the need for immediate human operation.

This study will address specific shortfalls identified by STMD, including intelligent multi-agent constellations, autonomy, edge computation, position, navigation, and timing for small spacecraft, small spacecraft propulsion, and space situational awareness (1625, 1438, 1433, 1557, 1431, 1430, 1589). The NASA Ames Mission Design Center (MDC) will provide subject matter expertise to support systems engineering trades, while experts in autonomy and spacecraft swarms in NASA's Intelligent Systems Division will lead the study and focus on identifying emerging next-generation swarm capabilities. The study objectives include: capturing the current state-of-the-art for multi-agent swarm capabilities, evaluating technologies and creating technology roadmaps, and developing at least one new technology demonstration mission concept. This initial NGS study report surveys the current state of the art in technology areas relevant for the next-generation multi-agent swarm design. Our primary focus is on surveying relevant deployed space systems[1], supplemented with selective analysis of relevant proposed missions and technology developments that have yet to fly.

For purposes of the NGS study, a *multi-agent swarm* is a multi-agent, coordinated, autonomous system distributed across multiple spacecraft. This definition builds on related terms[2]:

- *Autonomous spacecraft, autonomous system,* or *agent*: An entity able to make complex, important decisions independent from external control. Further, an agent's behavior can be modeled as a sense-think-act loop informed by the agent's goals and domain knowledge.
- *Multi-agent or autonomous multi-spacecraft system*: A system composed of multiple agents.
- *Coordinated autonomous system*: A multi-agent system whose decisions require interaction or consistency between agents. (As opposed to a constellation of independently operating agents that can solve fully decomposable system-level problems without interacting.)
- *Distributed autonomous system*: A coordinated autonomous system in which multiple agents participate in the system-level decision process or "think" part of the sense-think-act loop. (As opposed to using centralized control.)

These definitions place *multi-agent swarms*, the focus of the NGS study, at the center of the Venn diagram in *Figure 1*, where these various properties intersect. However, since very few multi-agent swarm missions have flown to date, our state-of-the-art survey will include technology developed for systems further out in the Venn diagram, in the space of all missions.

Traceability

The NGS study objectives trace to needs articulated by various NASA organizations. The study will address specific shortfalls identified in STMD's Civil Space Shortfalls (Lowry et al., 2024), including intelligent multi-agent

---

[1] We define a system as *deployed* if it was launched into space, whether or not the launch and subsequent operations were successful.
[2] These definitions are aligned with terminology in the field of multi-agent systems, although sources vary (Dorri et al., 2018; Stone & Veloso, 2000).

constellations (#1625), autonomy (#1438), edge computing (#1433), position, navigation, and timing (PNT) for small spacecraft (#1431), small spacecraft propulsion (#1430), and space situational awareness (#1589).

The NGS study will also address NASA Exploration Systems Development Mission Directorate (ESDMD) Technology Gaps from Appendix C of the Moon to Mars Architecture Definition Document (ADD) (Merancy et al., 2025). Relevant gaps include PNT on the Moon (#0101) and Mars (#0104) and reliable communication on Mars (#0102). Solutions for these gaps could include a communication and PNT (CPNT) swarm mission. Swarm missions could also validate autonomy technology relevant to integrated fault management (#1003), planning and decision-making (#1004), and autonomous monitoring (#1002).

NGS will also address technologies in the 2024 NASA Technology Taxonomy (Lowry & Charania, 2024). Nearly all of TX10 Autonomous Systems is relevant. There is a rough mapping between TX10 subtopics and our chapters:

TX10 Subtopic	Report Chapters
(not addressed)	Multi-Agent Communications
TX10.1 Situational and Self-Awareness	Multi-Agent Sensing
TX10.2 Reasoning and Acting	Multi-Agent Thinking, Multi-Agent Action
TX10.3 Collaboration and Interaction	Multi-Agent Operations
TX10.4 Engineering and Integrity	(cross-cutting)

Table 1: Taxonomy subtopics mapped to the chapters in this report.

Certain subtopics in other taxonomy sections are particularly relevant to multi-spacecraft systems, including TX4.2.6 Collaborative Mobility, TX4.4.2 Distributed Collaboration and Coordination, TX4.5 Autonomous Rendezvous and Docking, TX5.4 Network-Provided PNT, TX5.6 Orbital Debris Tracking and Management, and TX8.2.3 Distributed Aperture.

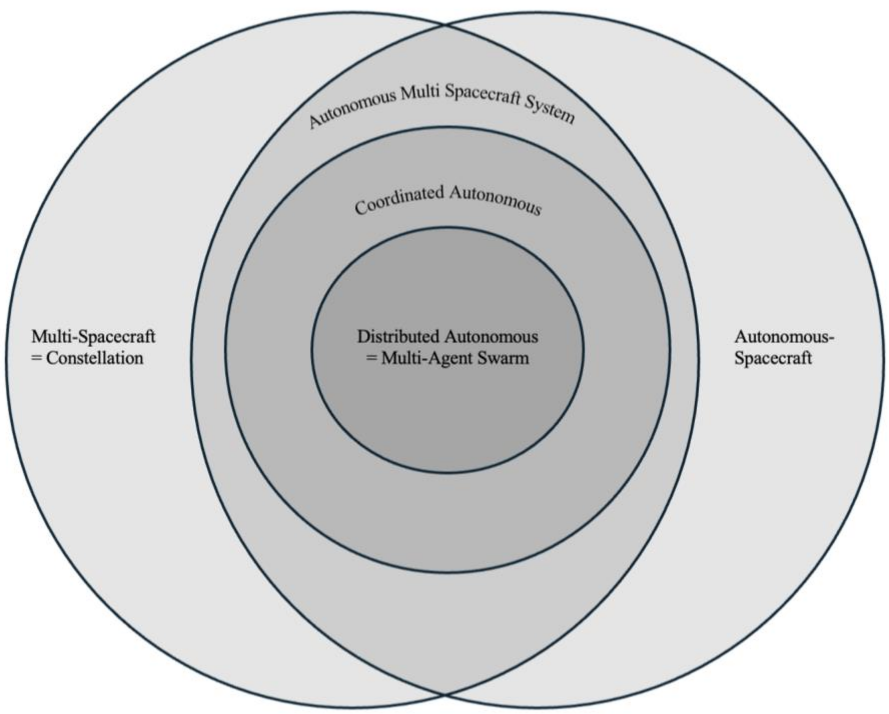

Figure 1: Multi-agent swarms are at the central intersection of various relevant system properties.

Related Work

This section positions our NGS SotA report in the context of the most relevant prior surveys. The NASA STMD 2024 State-of-the-Art of Small Spacecraft Technology report (Yost & Weston, 2024) focuses on complete individual spacecraft buses and subsystems of spacecraft, whereas our report focuses specifically on multi-spacecraft systems. Many topics in Yost and Weston's report are critical for enabling higher-order autonomy. While all core single-spacecraft technologies are needed to build a complete space system for autonomy demonstrations, specific technologies stand out as those which enable distributed systems and higher-order autonomy. Examples include complete market solutions for spacecraft buses, sensing systems for GNC, avionics systems, flight software, communications systems, and ground and data support systems. This report will discuss relevant subsystem technologies as they relate to multi-spacecraft systems.

The system property definitions used in this report expand on those used in (Adams et al., 2023), as described in *Table* 2 below. The table also compares our definitions to those used by the NASA Earth Science Technology Office (ESTO) Advanced Information Systems Technology (AIST) Program, in their New Observing Strategies (NOS), which defines a Distributed Spacecraft Mission (DSM) as "a mission that involves multiple spacecraft to achieve one or more common goals" (Le Moigne, 2023). In this study, any Multi-Spacecraft System (MSS) and NOS DSM would be equivalent. In contrast to traditional monolithic missions, where a single spacecraft incrementally achieves mission objectives, MSS/DSM both rely on a system of multiple spacecraft to accomplish mission objectives. These missions have become more prevalent due to the decreased launch costs and development costs per vehicle. MSS/DSM can involve heterogeneous spacecraft, consisting of different types, or homogeneous spacecraft, of all the same type.

However, it is important to note that an MSS does not necessarily imply distribution or autonomy. For example, the Concept of Operations (ConOps) for an MSS may involve manual commanding of individual spacecraft by ground operators. Additionally, some MSS examples are neither distributed nor autonomous, such as NASA's planned HelioSwarm mission (Spence et al., 2022). For HelioSwarm, operators will pre-compute plans for all spacecraft, which are then sent to a central hub spacecraft. The hub spacecraft relays execution plans to the individual detector spacecraft. Here, decision authority remains centralized, commanded by the ground segment and relayed to the central controller.

Other examples of MSS include SpaceX Starlink (Chaudhry & Yanikomeroglu, 2021a), Planet Labs' Flock (Boshuizen et al., 2014), NASA's CYGNSS mission (Carreno-Luengo et al., 2021), and NASA's TROPICS mission (Blackwell et al., 2018). The AIST NOS defines a *constellation* as "a space mission that, beginning with its inception, is composed of two or more spacecraft that are placed into specific orbit(s) for the purpose of serving a common objective," explicitly naming CYGNSS and TROPICS as examples. In our work, we also consider this definition of a constellation an MSS.

When autonomy is added to an MSS, we have Autonomous Multi-Spacecraft Systems (A-MSS). An A-MSS allows one of the spacecraft within the system to assume decision authority. Adding autonomy to a MSS offers various advantages, including addressing latency, bandwidth constraints, mission complexity, and transient or temporally limited events of interest. Moreover, autonomy enables the operation of spacecraft as a collective rather than as individuals. The Reconfiguration and Orbit Maintenance Experiments Onboard (ROMEO) experiment phase of the Starling 1.0 mission (Miller et al., n.d.) is one such example. The benefits of autonomy in A-MSS scale with the size of the system, emphasizing the significance of autonomy in larger A-MSS configurations.

NGS Study Definitions	(Adams et al., 2023) Definitions	AIST NOS Definitions
All Missions	= All Missions	= N/A / All Missions
Multi-Spacecraft System (MSS)	= Multi-Spacecraft System (MSS)	= DSM / Constellation
Autonomous-Spacecraft	Autonomous Multi-Spacecraft System (A-MSS) = the intersection of Autonomous-Spacecraft and MSS	A-MSS = ICC / Sensor Web
Coordinated-Autonomous	Subset of A-MSS	Subset of ICC / Sensor Web
Distributed-Autonomous	Subset of A-MSS	Subset of ICC / Sensor Web
Multi-Agent Swarm (Swarm or MAS)	Distributed Space System (DSS)	Subset of ICC / Sensor Web

Table 2: Equivalence of definitions commonly used in this report and other reports.

NASA's Surface Topography and Vegetation (STV) Incubation Study (Donnellan, 2021) identified several needs and gaps within the 2017 Earth Science Decadal Survey. The STV incubation study focuses on targeted observables, while the NGS team is observable-agnostic and focused on technology development. However, in the process of identifying Earth Science Decadal needs and gaps, several technologies relevant to this NGS report were identified. The STV study identified that "an architecture of multiple platforms and sensors on orbital and suborbital assets would address STV needs more thoroughly than a single orbiter," indicating a desire for a multi-spacecraft system. Furthermore, the authors state that "algorithm development and orbital constellation design will enable an STV architecture" and show that STMD and SMD have complementary technology development needs for future multi-spacecraft systems. Additionally, the STV incubation study, published in 2021, identified the following categories on their 3-7 and 7-10 year development roadmaps: *Sensor webs*, *new observation strategies design studies*, *architecture trade studies*, *onboard processing algorithms*, and *smart tasking*. This NGS study is therefore aligned with the selected STV categories on both topic and timeline, though the technology proposed is agnostic to observables. STMD requires that the NGS study produce at least one mission concept design, which may result in mission architectures relevant to STV. This places all overlapping STV and NGS needs within the same 2024 – 2031 window.

STV's category of *smart tasking* is of particular interest to this study, as it focuses on the use of constellations (multi-spacecraft systems) for intelligent observations of low latency, discrete events (such as landslides, earthquakes, volcanic eruptions, flooding, etc.). Distributed space systems are particularly useful for dynamic spatiotemporal phenomena where a human in-the-loop controller could not react in adequate time (due to network latency, distance, data limitations, etc.) (Adams et al., 2023). Additionally, STV describes the need to design a future mission with multiple spacecraft that leverages spacecraft information systems, multi-spacecraft coordination, and multi-sensor data fusion, calling for "technologies specifically designed for multi-spacecraft [systems] … through distributed processing onboard multiple spacecraft."

Distributed space missions within the STV study are identified as needing autonomous position, navigation, and timing as well as optimal geometric station keeping. These align with NGS study objectives for PNT and lunar orbital maintenance. Such technologies are critical for synchronizing events and sense-think-act sequences across multi-spacecraft systems. The STV incubation study expresses the need to "operate as a sensor web with nodes that are interconnected by a communications fabric and that function as a single, highly coordinated, virtual instrument." This need involves all NGS relevant domains; details are broken down within the report chapters.

Finally, the Autonomous Navigation Demonstration Relevance Assessment Team (ANDRAT) conducted a study on autonomous onboard navigation for future space-based missions that cannot rely on the Global Positioning System (GPS) (*Andrat-Final-Report.Pdf*, n.d.). The ANDRAT developed a framework to assess and communicate the maturity of autonomous navigation systems, in the form of levels of autonomy, shown in *Table 3*. At each level, the role of the onboard decision-maker and ground support are described. The rightmost column maps prior missions into these levels for context. Based on these definitions, ANDRAT recommended that a technology demonstration mission target Level 3 autonomy, as it enables the onboard navigation system to make simple decisions over multiple mission phases while also relying on a ground support team to make more complex decisions. This autonomy level is highlighted in light blue in *Table 3* for clarity. Overall, the report compiled by the ANDRAT highlights critical gaps that must be closed for higher levels of autonomous navigation and provides a concrete roadmap for how future missions can advance the state of the art.

	Level	Spacecraft role	Ground operator role	Example missions
0	No automation		Process/analyze data, make decisions, upload to spacecraft	
1	Low-level automation	Prescribed (e.g., rule-based) low-level processing and analysis	Make decisions using information from low-level analyses on spacecraft; independently verify	Stardust, Lucy
2	Brief, limited automation	Automated processing, analysis, and prescribed decisions during specific/single phases of the mission within the expected operating envelope	Perform all functions during all other phases and decide if/when to enter automation	DS1, Deep Impact, OSIRIS-REx, DART
3	Extended, limited automation	Automated processing, analysis, and prescribed decisions across most of the mission within the expected operating envelope	Outside of expected operating envelope, intervene with tasks or updated software parameters	DS1 (extended duration, but single phase)
4	Supervised autonomy	Processing and analysis, recommendations on complex decisions, explanations conveyed even outside of the expected operating envelope	Review and independently verify recommendations before sending decisions. Operator can intervene/override when needed (e.g., major faults).	
5	Full autonomy	Processing, analysis, and complex decision-making even outside of the expected operating envelope; onboard verification of decisions	Update high-level objectives or target orbits. Intervention from ground is only necessary when failure is expected.	
6	Intelligent autonomy	Processing, analysis, and complex decision-making even outside of the expected operating envelope; onboard verification while learning and adapting	Update high-level objectives or target orbits. Intervention from ground is only necessary when failure is expected.	

Table 3: ANDRAT navigation autonomy levels for spacecraft.

Advantages of Swarms

We find the ESTO AIST NOS categorizations more than adequate to describe Earth observation systems, but our classifications are more useful when defining the scope and motivation of the NGS's contributions to boarder space systems research. In particular, the distribution of decision authority in a Multi-Agent Swarm has several advantages over the centralized authority in a constellation and is motivated by a number of factors, included but not limited to the following:

- Fault tolerance and redundancy: By distributing decision-making authority among multiple spacecraft, the system becomes more resilient to individual failures.
- Improved scalability: Distributed decision making enables the system to remain operational when adding and removing agents dynamically.
- Increased computational capacity: Distributing decision-making among multiple spacecraft allows for parallel processing and sharing computational load.
- Enhanced adaptability and flexibility: Distributed decision-making enables spacecraft to make autonomous and localized decisions based on local perception.
- Efficient task allocation and coordination: By distributing decision-making, spacecraft can autonomously allocate tasks among themselves based on their capabilities, proximity, and availability.
- Increased robustness to communication delays and failures: In distributed decision-making systems, spacecraft can continue operating even in scenarios where communication between nodes is intermittent or disrupted.

The utilization of MSS has become increasingly prevalent in space exploration. As such, it is essential to determine when a swarm is the most suitable approach for a given mission. A swarm is not necessary in situations where data can be transmitted back to Earth at a slow rate without compromising mission objectives, or where increasing onboard storage capacity is a viable solution. Additionally, if a single spacecraft can provide adequate coverage, or if additional coverage can be achieved through other means, such as ground-based assets, a swarm may not be required. Furthermore, if human operators can effectively control the system in real-time, or if the optimal solution or mission requirements are relatively stable and do not change frequently, a simpler system may suffice.

Compelling Swarm Scenarios

A swarm is ideal for missions requiring rapid response, such as low-latency delivery of data or rapid follow-up observations of fleeting and unpredictable phenomena. These systems also enable high-resolution spatiotemporal coverage, making them suitable for missions requiring detailed observations of dynamic events or phenomena over large areas. Moreover, in environments with high communications latency due to distance or cost, a swarm can operate effectively in a decentralized manner, reducing reliance on continuous communication with Earth. In short, if any one of the following are true, then a swarm may be an ideal solution:

1. Missions requiring rapid response low-latency delivery of data or rapid follow-up observations of fleeting and unpredictable phenomena
2. Missions requiring high-resolution spatiotemporal coverage
3. Missions in which communications latencies are high due to distance or cost

Evaluating Swarm Technologies

To evaluate the current state of the art in multi-agent swarm systems, this report describes metrics to assess technologies within five independent domains: **Multi-Agent Communications**, **Multi-Agent Sensing**, **Multi-Agent Thinking**, **Multi-Agent Action**, and **Multi-Agent Operations**. Each domain includes missions that demonstrate domain-specific capabilities, and missions are evaluated according to domain-specific metrics. The five domains are illustrated on a radar chart in *Figure 2*. Each axis corresponds to a domain, and moving further out radially indicates a "more capable" system within that domain. Our approach is to assess the performance and maturity of missions within each domain using a shared lens. In addition to this, we also evaluate the top-performing missions in each domain across all axes simultaneously (see *Conclusion*). This cross-domain evaluation illustrates the "area" these missions occupy – some are extremely capable in multiple domains, while others are more one dimensional. When plotted on a radar chart, the most capable mission would cover the most area.

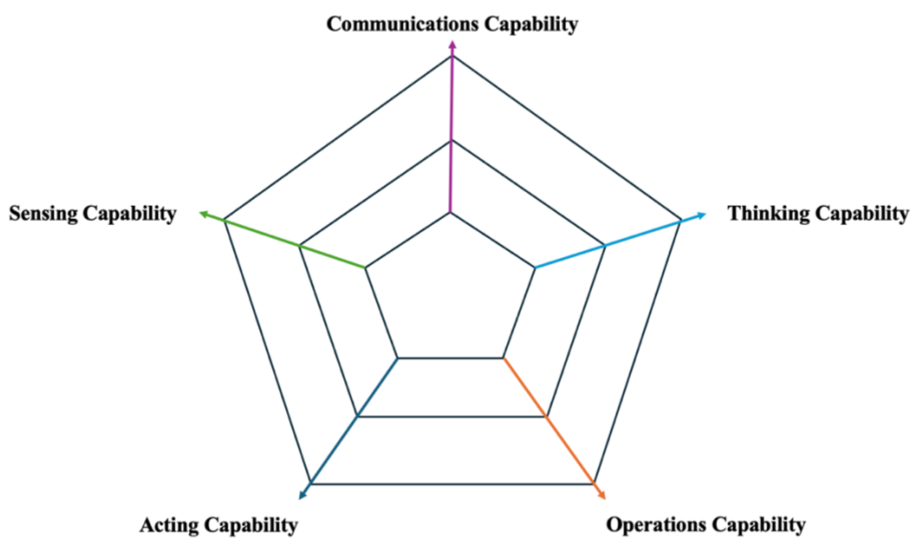

Figure 2: A radar chart of the five domains in multi-agent systems. Moving farther out radially corresponds to a more capable system.

We now describe each domain in more detail. **Communications Capability** assesses how agents exchange information, considering factors such as network structure, latency, bandwidth usage, and availability. **Sensing Capability** evaluates the accuracy, speed, resolution, and coverage of the swarm's sensing, perception, and awareness capabilities. It captures how well the system perceived its environment through distributed sensing across the swarm. **Thinking Capability** reflects the collaborative decision making of the system, including how decisions are made in a distributed or decentralized manner, temporal responsiveness, planning depth, and optimality. **Acting Capability** measures the degree of coordination, execution accuracy, and responsiveness required for agents to interact with the environment and one another effectively. Finally, **Operations Capability** reflects the management of the swarm as a whole, including how it is commanded, observed, and maintained throughout the lifetime of a mission.

In addition to domain-specific metrics, we incorporate cross-cutting metrics that apply across all five domains. One of these is **scalability**, which assesses the number of agents or nodes in a multi-agent system. As swarm size increases, challenges related to network structure, decision-making authority, coordination, and overall manageability often increase. Evaluating how system performance scales with the number of agents provides valuable insight into how robust a particular architecture is and its potential for operational deployment at larger scales.

Another example of a cross-cutting metric we introduce is the **swarm score**, which reflects how "swarm-like" a mission is. The swarm score is computed by determining how many relevant system properties a mission intersects in the Venn diagram in *Figure 1*. The score can range from 1-4, where 1 indicates a mission with minimal swarm characteristics and 4 indicates a mission with all key swarm characteristics. Swarm characteristics include coordination, distribution, and level of autonomy, as illustrated in *Figure 1*. For example, a constellation of independently operating spacecraft, such as COSMIC-2 (Weiss et al., 2022), would receive a swarm score of 1, while a multi-agent, coordinated, autonomous system distributed across multiple spacecraft, such as the StarFOX and DSA software payloads flown on NASA's Starling mission (Kruger et al., 2023; Miller et al., n.d.), would receive a swarm score of 4. This metric provides a way to assess swarm maturity, capturing key swarm characteristics, and enabling a clear differentiation between systems that are multi-agent in nature versus those that demonstrate true swarm-like behaviors.

Together, the domains **Multi-Agent Communications**, **Multi-Agent Sensing**, **Multi-Agent Thinking**, **Multi-Agent Action**, and **Mult-Agent Operations** provide a comprehensive framework for evaluating current and emerging multi-agent systems. When quantitative data is available for a domain-specific metric (for example, communications bandwidth or data resolution), it supports the assessment of missions in that domain. However, each domain is characterized by multiple metrics, which may include both quantitative and qualitative measures. To ensure the applicability of these metrics across all missions in a domain, this report combines them in a qualitative way. This allows for meaningful comparisons across systems with varying data availability, levels of maturity, and operational goals. Additionally, it supports the identification of key strengths, trade-offs, and gaps in each domain that can guide future advancements in multi-agent swarm systems.

Multi-Agent Systems Domains

The remaining chapter structure of this report follows the domains. Here are the domains, their scopes, and definitions:

Multi-Agent Communications capabilities enable sending and receiving data between agents in the multi-agent system. Example themes include:

- Crosslink, where agents transmit and receive data directly with one another, as opposed to requiring the transmission be relayed through a ground station
- Quality of service (QoS) protocols that provide properties like guaranteed delivery, in-order delivery, or efficient large file transfer on top of unreliable packet delivery
- Secure communication protocols

Multi-Agent Sensing capabilities enable system-level perception of the environment through distributed data collection. Example themes include:

- Analyzing in situ measurements with a formation of spacecraft to map processes that vary across space and time (e.g., space weather)
- Analyzing remote sensing observations from multiple spacecraft that provide a stereo baseline or synthetic aperture, including data association for object tracking
- Calculating joint position, navigation, and time estimates from signal timing data collected on crosslinks and space-to-ground links by multiple spacecraft

Multi-Agent Thinking capabilities enable collaborative decision-making to select coordinated actions or plans that achieve system-level goals. Example themes include:

- Collective mission planning and resource allocation
- Multi-agent planning and scheduling
- Joint knowledge and understanding
- Behavior and intent prediction
- Multi-agent fault management

Multi-Agent Action capabilities enable agents to jointly execute their selected coordinated actions or plans. Example themes include:

- Formation flying, where agents perform coordinated maneuvers into desired spatial configurations to accomplish mission objectives
- Motion control and trajectory following, enabling agents to maintain these formations by adjusting thrust and attitude
- Fault recovery and response, allowing agents to autonomously reconfigure, redistribute tasks, or adapt actions to ensure mission success in the presence of failure

Multi-Agent Operations capabilities enable one-to-many user control of a multi-agent space system, expanding the context to include ground control. Example themes include:

- Adaptive mission management interfaces for operators
- Telemetry summarization and visualization for multi-agent systems
- Scalable ground system architectures for multi-agent system management

State of the Art – Existing Multi-Agent Systems

The landscape of multi-agent swarms spans a wide range of mission sizes, objectives, and technology readiness levels. The NGS study team assembled as catalog of 50 known multi-agent systems, including a few single-agent systems when deemed relevant. These missions range from small CubeSat technology demonstrations with just two nodes to massive operational constellations numbering in the thousands. The catalog captures past, current, and future multi-agent space-based systems around Earth and the Moon, including government and industry science missions, commercial internet constellations, and technology demonstrations. For systems that were fully deployed, we report the peak node counts, while for planned missions, we include the projected number of maximum nodes. The complete list of systems, including individual spacecraft characteristics and operational details, is provided in the *Appendix*.

The catalog reveals several patterns in the current state of the art. The smallest missions, such as the University of Toronto's CanX-4 and CanX-5, served primarily as a technology demonstration with just two nanosatellites, focusing on precision formation flying and inter-satellite communications. In general, these types of missions were launched into low Earth orbit (LEO), had masses < 10 kg per node, and had lifetimes < 1 year, demonstrating that even the smallest missions can push the state of the art. On the other end of the spectrum, commercial systems like SpaceX's Starlink represent the largest multi-agent systems ever deployed, with thousands of nodes in orbit. These constellations involve little to no coordination between agents but demonstrate robust mesh networking and long-term operations at scale.

Between these extremes lies a diverse collection of missions. ESA's Cluster II and NASA's MMS each flew four identical nodes to probe the magnetosphere, while mission planning for the A-Train constellation coordinated heterogeneous platforms to enable cross-calibrated Earth observations. Planned missions like HelioSwarm extend this concept by distributing sensing across many small spacecraft with different roles. Taken together, several key features emerge. Many of the most capable missions were operated as swarms, such as NASA's StarFOX and DSA software payloads, which demonstrated coordinated autonomy and distributed decision-making across multiple spacecraft. Scientific missions, such as COSMIC-2 and CYGNSS, highlight distributed high-frequency, high-precision observation from LEO. Operational concepts vary widely: some systems rely on one-to-many command links and others employ peer-to-peer crosslinks with partial or full mesh topologies. However, no flown multi-agent system has successfully demonstrated divide-and-conquer planning strategies, high-fidelity sensor fusion, or coordinated and distributed fault recovery. Further, almost all known multi-agent missions operated around Earth, with the advantage of GPS for precise PNT. To extend multi-agent capabilities to other environments, more sophisticated and robust technologies are needed.

Multi-Agent Communications

Introduction

Multi-Agent Communication is concerned with how each agent exchanges information with other agents in the system. This domain was selected because it describes how agents in the system coordinate with one another and act as a unified swarm, rather than act as a constellation of independently operating agents. This capability is foundational to all other domains. As such, multi-agent communication is fundamental to the study of the next generation of multi-agent swarms. This section is only concerned with agents that directly communicate with one another. Multi-Agent Communications is not concerned with infrastructures such as the reverse bent pipe model, in which each agent communicates only with the ground.

Key Capabilities

- Crosslink
- Quality of Service (QoS)
- Interoperability

Crosslink is the capability for agents to have direct communication with one another. It enables agents to coordinate with each other without the need to interact with the ground. Crosslink is fundamental to a next generation multi-agent swarm, as it is not always practical or possible for a swarm to have to first communicate with the ground before be performing actions. In some cases, requiring the swarm to communicate with the ground would lead an operations delay that could cause the swarm to miss out on quickly transient events. This is especially true for swarms working in deep space where the propagation delay is very large.

The next capability is **Quality of Service (QOS)**. According to IEEE 802.1DC, Quality of Service (QoS) functions are those that affect the network's ability to serve data and are measured by latency, frame loss probability, and the variability of the two ("IEEE Standard for Local and Metropolitan Area Networks—Quality of Service Provision by Network Systems," 2024). Latency measures the time required to transmit a data packet from source to destination. Frame loss probability measures the likelihood that a data packet will not make it to its destination. In an ideal multi-agent swarm, latency and frame loss are kept as low as possible. However, operating in a space environment presents several challenges that lead to higher latency and frame loss. A next generation multi-agent swarm must incorporate a quality of service capability that enables communications that are robust against high latency or frame loss.

As the commercial space industry grows, so will the number of disparate space systems. As such, collisions between these systems will become more probable. Currently, collision avoidance begins with the ground identifying a possible collision. But as the more and more systems become operational, timely responses will become more critical, and these disparate systems would benefit from having direct communication with one another. A next-generation swarm should have a communications system that incorporates **interoperability** with other systems. Interoperability is the capability of a system to interact with different systems and exchange information. In the context of multi-agent swarms, adopting an interoperable standard would allow satellites from different commercial providers to not only warn each other of a possible collision in a timely manner but also negotiate who will make a collision avoidance maneuver.

Communications Capability Metrics

This section captures metrics that are used to assess how advanced a multi-agent communications system is with regards to each of the capabilities identified above. Examples are provided for what it means for a multi-agent communications system to be simple versus advanced with respect to each capability. These examples are not intended to be an absolute scale. Rather, they provide a general sense for what it means for a system to be simple versus advanced.

- Routing Efficiency

- QoS negotiation
- Open Communications Standard

Routing efficiency was chosen as a measure of how advanced a multi-agent swarm's crosslink capability is. Routing efficiency refers to how the network chooses to route packets to reduce latency and frame loss. In a simple multi-agent swarm, the mesh networking protocol can provide a basic routing capability. Routing in a simple case may be static in the sense that the topology cannot be reconfigured over the life of the mission. There are very few optimizations in a simple multi-agent swarm. However, in a more capable multi-agent swarm system, routing would be dynamic and optimize the network topology to minimize latency and frame loss. It would not only seek the path with the fewest hops to the destination, but choose a path that optimizes stability, latency, and throughput. This capability would require additional layers of logic as the swarm would have to monitor the state of the network and optimally decide how to route their data from source to destination.

QoS negotiation was chosen as a measure for how advanced a multi-agent swarm's quality of service capability is. Quality of Service is measured by latency and frame loss. How a multi-agent system negotiates between these two measurements is referred to as QoS negotiation. A simple communications system for a next generation multi-agent mission would provide, at minimum, fixed set of quality of service parameters to ensure that the swarm can meet its goals. These are "fixed" in the sense they would be defined on the ground and not runtime configurable based on the current configuration of the multi-agent swarm. Engineers would define quality of service parameters to best fit the overall mission profile. This approach is simple, but it may not be optimal in some phases of the mission. A more capable multi-agent communications system would enable the next generation multi-agent swarm to negotiate the quality of service based on the current state of the swarm. This would require the agents to analyze their needs and capabilities and negotiate the best possible quality of service possible. This negotiation adds an extra layer of logic for the sake of optimality.

For this assessing the capability of a multi-agent swarm's communication system with regards to interoperability, the metric chosen was whether or not the system adopted an **open communications standard**. An open standard is a standard that can be utilized by anyone. A less capable system would adopt a closed communication standard that is either wholly owned by a single organization or is only shared among a few organization. A closed standard limits the ability of a swarm to communicate with other swarms from. Conversely, a more capable system would adopt an open standard that would allow their swarm to exchange data with agents from other swarms.

State of the Art – Complete Swarm Systems

This section reviews the current state of the art of multi-agent communications in space systems, with a focus on flown missions involving more than one spacecraft. Key capabilities include routing efficiency, QoS negotiation, and communications security. Standout missions include NASA's Starling and SpaceX's Starlink, which have successfully demonstrated multi-agent mesh networking in orbit, with 4 nodes to as many as 7,000 nodes, respectively. These missions represent the state of the art in multi-agent communications and showcase the scalability, resilience, and adaptability of distributed networks in space. The overall Communications Capability of each missions is qualitatively evaluated using the metrics identified earlier. Because of the scarceness of swarm missions, this study expands its scope to include notable future missions and notable swarm missions that were not fully successful.

The **CubeSat Laser Infrared CrosslinK (CLICK)** mission (Mayer et al., 2019) is a technology demonstration of laser communication between two CubeSats. CLICK was chosen as a representative mission for multi-agent communications because it demonstrates crosslink capabilities. A total of 3 spacecraft were built for CLICK: CLICK A, B, and C. CLICK A was launched in 2022 and flew alone as a risk reduction mission to test the laser hardware. CLICK B and C are scheduled to fly no earlier than the second quarter of 2026. In terms of multi-agent communication, CLICK ranks lowest on the capability scale, as communication will only take place between two nodes. Because CLICK consists of only two nodes, routing between CLICK B and C will be static, and therefore it would score very low in routing efficiency, as there is no way optimize the network topology.

HelioSwarm is a multi-satellite science mission to characterize the solar atmosphere (Spence et al., 2022). HelioSwarm consists of a single hub satellite and nine smaller satellites. HelioSwarm has a static hub and spoke or a star network topology. All communications between the HelioSwarm nodes spacecraft and the ground are routed through the larger hub spacecraft. HelioSwarm scores lower in communications capability because the networking topology is static and cannot be modified to optimize routing between nodes. Notably, Helioswarm nodes will operate at different distances from the hub in distinct mission configurations governed by the science, so even though the network topology remains unchanged, the network must operate seamlessly given under a range of distance and communication angle constraints.

The **PY4** mission is a low-cost, rapid-turn-around, four satellite technology demonstration based on the PyCubed avionics framework (Holliday et al., 2024). PY4 was chosen as a representative mission for multi-agent communication because it demonstrates the crosslink capability. PY4 is designed with fault-tolerant ad-hoc mesh network that supports the ability to store-and-forward ground commands (Stupl & Holliday, 2023a). There is little information detailing how the fault-tolerant ad-hoc mesh network was achieved. However, fault-tolerant mesh networking would require PY4 to reconfigure the network topology to efficiently route around instability. PY4 is considered to score higher on the communications capability scale because it demonstrated a dynamic mesh network between four nodes.

The Edison Demonstration of SmallSat Networks (EDSN) mission was a project to demonstrate satellite networking. EDSN built a total of 14 CubeSats for the mission, 8 of which would fly on Super Strypi, a modified sounding rocket. However, due to a launch failure, all CubeSats were destroyed. Two of the remaining nodes were modified to be deployed from the ISS in a follow-on mission called **Network & Operation Demonstration Satellite (NODES)** (Hanson et al., 2014). EDSN and NODES were selected as representative missions for multi-agent communications because they demonstrate crosslink between spacecraft. EDSN/NODES rank higher on the communications capability scale, as they were designed to autonomously determine which spacecraft is best suited to control the space network. This would have demonstrated a dynamic network topology that can reconfigure itself to optimize routing efficiency. EDSN/NODES was placed above CLICK because, despite the launch failures that led to only 2 spacecraft flying, EDSN/NODES was designed with the capability to have a dynamic network topology that could be autonomously reconfigured across several nodes.

The **Starling 1.0** mission (Miller et al., 2024), led by NASA Ames Research Center, is a satellite swarm of four 6U CubeSats launched July 2023. Starling demonstrated many firsts in spacecraft swarm autonomy. With regards to multi-agent communication, Starling utilized the Better Approach to Mobile Ad-hoc Networking (B.A.T.M.A.N) Advanced protocol to manage and route crosslink traffic between the four spacecraft. B.A.T.M.A.N Advanced provides an abstraction layer over a computer's network devices that allows packets to be routed around unstable connections (*B.A.T.M.A.N. Advanced — Open-Mesh.Org Latest Documentation*, n.d.). The Starling crosslink radios were treated as just another network interface as far as the flight computer was concerned. As a result, Starling could dynamically route packets more efficiently through the network. This capability allows Starling to score much higher on the routing efficiency metric.

The SpaceX **Starlink** constellation provides satellite internet services to customers in regions where traditional internet service is not available. The Starlink scores very high in the communications domain due to several factors. The first is the sheer number of satellites that are communicating with one another. Starlink states they have more than 7,800 satellites in orbit[3] that have crosslink capabilities (Chaudhry & Yanikomeroglu, 2021b). The next factor is how the Starlink is able to decide how to efficiently route connections through the network. Though there is little literature published regarding how SpaceX is able to route connections between satellites, the use of laser communication means that the Starlink satellites must determine which neighboring satellite is best suited to forward the connection, demonstrating routing efficiency.

Summary

A next generation multi-agent swarm requires a multi-agent communications system that can meet the dynamic needs of the swarm. This includes capabilities such as crosslink to provide a networking infrastructure, quality of

[3] https://www.starlink.com/updates/network-update

service to ensure data is reliably transferred, and interoperability to enable communication between disparate swarms.

Figure 3 provides an overview of how each system compares. At the highest level of capability is SpaceX's Starlink satellites. Although not technically a swarm, Starlink was given the highest capability rating due to the sheer number of spacecraft and the use of laser communication between them. NASA's Starling was evaluated at the second highest, as it uses a dynamic mesh network that can efficiently route connections around network instability. Next was EDSN/NODES, which could optimize network routing by deciding which of the spacecraft were the best for communicating with the ground. PY4 sits nearly alongside EDSN/NODES as it also demonstrates a mesh network that could reconfigure itself around faults. HelioSwarm sits lower on the graph because it utilizes a static hub and spoke or star topology that cannot be reconfigured. Lastly is CLICK, which demonstrates a static network between only two nodes.

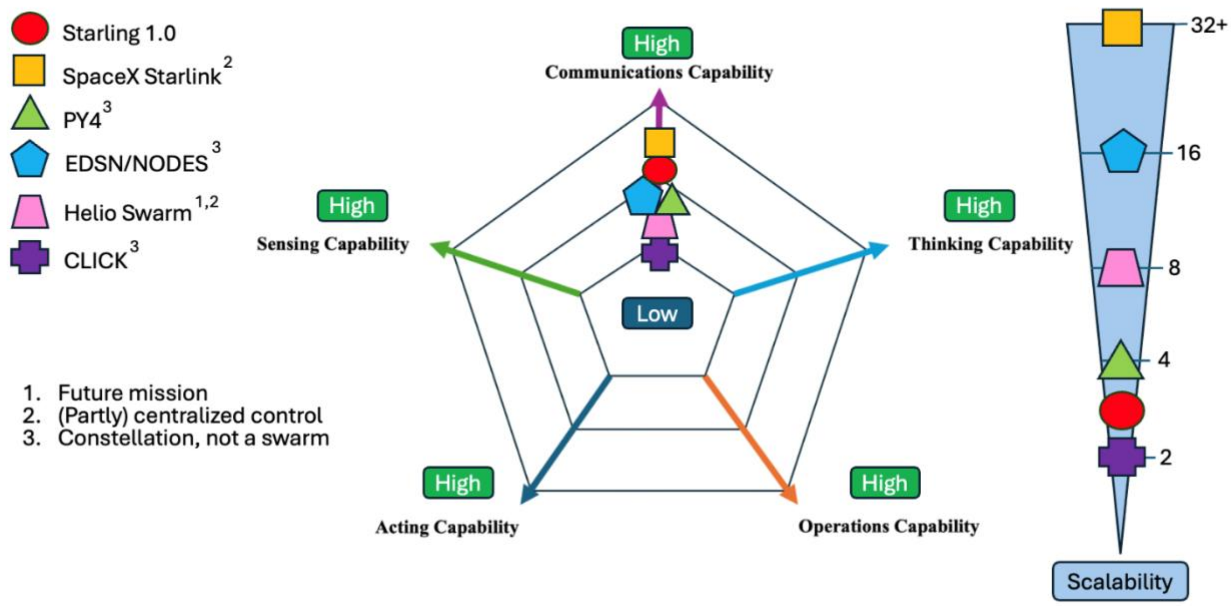

Figure 3: A radar chart of the five domains in multi-agent systems, with the state-of-the-art multi-agent communications missions plotted. Moving farther out radially corresponds to a more capable system.

Multi-Agent Sensing

Introduction

Multi-Agent Sensing refers to the collective ability of a system to perceive its environment through observations across multiple agents. This domain is foundational for shared situational awareness, decision-making, and safety. Rather than relying on a single platform, multi-agent sensing systems integrate data from multiple agents to achieve greater spatial and temporal coverage. In space-based swarms, specifically, agents often perceive their own states, those of neighboring agents, and external environment features under varying observational constraints.

Key Capabilities

- Data Collection
  - Observation
- Perception
  - Detection
  - Recognition
  - Tracking
  - Mapping
- Awareness
  - Attitude Estimation
  - Orbit Determination

Key capabilities in this domain begin with data collection, which encompasses the fundamental task of gathering raw information from the environment. This includes **observation**, where agents record information from various sensors, such as optical imagery, radar returns, or infrared signatures. Perception builds upon this by enabling agents to interpret their surroundings using the collected data. Perception includes **detection**, which enables agents to identify the presence of physical objects or features of interest in the observed data, such as a spacecraft, terrain obstacles, or anomalies. **Recognition** extends this by classifying detected objects or features into known categories. For example, identifying a resident space object as a particular spacecraft or distinguishing geological formations on a planetary surface. **Tracking** builds upon these capabilities by maintaining continuous estimates of the state of a target over time, such as its position, velocity, or trajectory. This allows agents to follow dynamic targets, predict their future behavior, and adapt sensing or navigation strategies accordingly. Finally, **mapping** combines distributed observations into a coherent representation of the environment, whether that is a topographic model, a hazard map, or a feature layer used for subsequent planning.

Another core sensing capability is awareness, which refers to the ability of an agent to infer knowledge about its own state, the state of other agents, and the surrounding environment. Key components of awareness include **attitude estimation**, which estimates the orientation of an agent relative to a defined reference frame, and **orbit determination**, which determines target orbital elements over time and space. These capabilities are especially crucial in deep space environments where external navigation aids, such as GPS, are unavailable.

Sensing Capability Metrics

- Accuracy / Precision
- Speed
- Resolution
- Coverage / Revisit Frequency

To evaluate the state of the art of Multi-Agent Sensing, this report introduces a set of metrics which capture the performance of sensing capabilities in multi-agent systems. These metrics are qualitatively combined to assess the Sensing Capability of relevant missions and technologies. **Accuracy** reflects how closely sensor outputs match true environmental states or features. High accuracy ensures that decisions and actions based on the sensor data are reliable. **Speed** encompasses sensing latency and responsiveness. Low-latency systems allow the swarm to react to fast-evolving situations in real-time. **Resolution** measures the spatial and temporal granularity of the sensed data. Finer resolution supports more informed decision making. **Coverage and revisit frequency** assess how broadly and how often a sensing system observes a given area or target. Wide coverage prevents blind spots, while high revisit rates enable timely updates. The balance between these metrics determines the overall Sensing Capability of a particular system.

Together, these capabilities and metrics provide a foundation for evaluating the state of the art in Multi-Agent Sensing. As sensing architectures grow more autonomous and distributed, the metrics help highlight the trade-offs, strengths, and technological gaps, ultimately guiding the development of more capable swarm systems. This chapter primarily focuses on flown multi-agent missions that have successfully demonstrated the identified sensing capabilities. Select examples from planned multi-agent missions or single-agent missions are also included when they illustrate capabilities relevant to a next generation multi-agent mission. In cases where no multi-agent system has flown yet for a given capability, the gaps are noted to highlight areas for future advancement.

State of the Art – Complete Swarm Systems

This section reviews the current state of the in multi-agent systems, concentrating on flown missions that involve more than one spacecraft. Key capabilities include data collection, perception, and awareness, as previously defined. Technologies such as the Starling Formation-Flying Optical Experiment (StarFOX) payload on the NASA Starling mission have demonstrated advanced levels of onboard perception, including autonomous detection, recognition, and tracking of unknown targets using multiple observers. StarFOX, and missions beyond Earth, such as the Gravity Recovery and Interior Laboratory (GRAIL), also demonstrate state-of-the-art multi-agent awareness capabilities. These missions represent significant progress in space situational awareness, showcasing how multiple spacecraft can collaborate to generate maps of their surroundings and navigate in GPS-denied environments.

Data Collection

Data collection is a foundational sensing capability that simply involves using a sensor to take measurements from the environment. An important element of this is observation – the collection of raw sensor data. This data provides the baseline information required for higher-level sensing tasks, such as detection, recognition, and mapping. The following section reviews the state of the art in data collection for multi-agent systems, focusing on mission technologies that have been demonstrated in flight. When relevant, planned multi-agent missions are discussed to highlight areas that are actively advancing. The overall Sensing Capability of each mission is qualitatively evaluated using the metrics identified earlier.

COSMIC-2, also known as FORMOSAT-7, is a constellation of six satellites operated jointly by the Taiwanese Space Agency and the National Oceanic and Atmospheric Administration (NOAA). The mission, launched in 2019, was designed to provide advances in meteorology, ionosphere climatology, and space weather research. Each satellite carries a Tri-GNSS Radio-Occultation System (TGRS) and Ion Velocity Meter (IVM) payload which produce high-resolution vertical profiles of atmospheric and ionospheric parameters, which are then processed on the ground. The volume and precision of these observations – spanning thousands of profiles per day with a median 30 minute latency – highlight the capability of COSMIC-2 to perform high-fidelity, low-latency perception of space weather. The TGRS instrument demonstrated acceptable accuracy in atmospheric and ionospheric measurements, while the IVM instrument achieved plasma density measurements with a precision of 1% or better (Weiss et al., 2022). These results put COSMIC-2 at the forefront of the state of the art for multi-agent observation in space.

Launched in 2007, **THEMIS** (Time History of Events and Macroscale Interactions during Substorms) was a NASA mission consisting of five identical microsatellites that studied the magnetospheric processes responsible for auroral sub-storms around Earth. In 2010, the two outer nodes were redirected into a lunar orbit, where they became the **ARTEMIS** (Acceleration, Reconnection, Turbulence, and Electrodynamics of the Moon's Interaction with the Sun) mission. The goal of ARTEMIS was to collect simultaneous measurements of energetic particles around the Moon to better understand how the solar wind interacts with the lunar surface and magnetic environment. These multipoint data measurements were used on the ground to build 3D maps of the Moon's magnetic field and revealed new information on the structure of the lunar wake (Angelopoulos, 2008; Bester et al., 2014; *THEMIS (Time History of Events and Macroscale Interactions during Substorms) - eoPortal*, n.d.; *THEMIS-ARTEMIS - NASA Science*, 2017).

As a constellation repurposed across Earth and lunar orbits, THEMIS-ARTEMIS exemplifies the state of the art in multi-agent perception, demonstrating how sensors distributed across multiple spacecraft can jointly resolve complex plasma interactions that would be impossible for a single spacecraft to do alone.

The **Bright Target Explorer (BRITE) Constellation** is a multi-national astrophysics mission that demonstrated successful multi-agent observation through coordinated monitoring of the brightest stars in the sky. BRITE comprises of a constellation of 6 nanosatellites launched between 2013 and 2017 by Austria, Canada, and Poland. The mission was specifically designed to perform space-based photometry of hot, massive stars. Each spacecraft is equipped with a small optical telescope and a CCD detector, enabling it to capture images of stellar targets with high photometric accuracy and temporal resolution. Specifically, the photometric brightness measurements collected by BRITE achieved a root-mean-square scatter of about 0.001 magnitudes or better, during a typical 15 minute observation window (Handler et al., 2017; Pablo et al., 2016; Popowicz et al., 2017). This data is then processed on the ground to produce precision light curves that correct for instrumental effects and noise. By distributing observations across multiple satellites with slightly offset orbits and complementary viewing geometries, BRITE achieves continuous monitoring of selected stars – capabilities that are difficult or impossible to achieve from ground-based observatories due to atmospheric interference and limited visibility windows. This multi-point architecture enhances spatial and temporal coverage, enabling the detection of subtle stellar variability. As a result, BRITE demonstrates how coordinated nanosatellite systems can deliver scientifically valuable astrophysical observations.

Similarly, the **European Space Agency (ESA) Swarm** mission, launched in 2013, consists of a constellation of three satellites with the objective of measuring the magnetic field of Earth and near-Earth electromagnetic environments. The main instruments on board each satellite are a vector and a scalar magnetometer, which enable Swarm to make multi-point measurements of the strength, direction, and variation of magnetic fields. When combined on the ground, this data provides the necessary observations required to model various sources of the geomagnetic field of Earth and its interaction with other physical aspects of the Earth system. The ESA Swarm mission achieved high precision in its magnetic field measurements. The in-situ magnitude of the magnetic field is measured with a random error below 0.3 nanoteslas (nT) at a 2σ confidence level, and the magnitude remains stable within 0.12 nT over a three-month period. Vector components of the magnetic field are determined with random errors under 1 nT, while their stability is maintained at better than 1 nT per year (Friis-Christensen et al., 2008; Iorfida et al., 2023; *Swarm (Geomagnetic LEO Constellation) - eoPortal*, n.d.). These performance metrics underscore the capability of the ESA Swarm mission at delivering high-fidelity, reliable data critical for detailed characterization of the geomagnetic environment of Earth – a feat that is achieved due to the multi-agent nature of the mission.

The **Polar Radiant Energy in the Far-InfraRed Experiment (PREFIRE)** is a NASA mission consisting of two CubeSats launched in May 2024 to study the polar regions of Earth. The spacecraft were placed in asynchronous near-polar orbits to enable collocated measurements of the same locations separated by 0-12 hours. Each satellite carries a thermal infrared spectrometer designed to measure far-infrared radiation, which dominates the energy emitted from the poles yet has been poorly observed. By operating as a distributed pair, the PREFIRE satellites capture the temporal variability of rapidly evolving processes, such as cloud formation, sea ice changes, and surface temperature dynamics (L'Ecuyer et al., 2021; *PREFIRE - NASA Science*, 2023). This coordinated sensing approach provides a precise view of polar energy dynamics that a single satellite could not capture alone, making PREFIRE a state-of-the-art mission in multi-agent perception.

Launched in March 2025, the **Polarimeter to UNify the Corona and Heliosphere (PUNCH)** mission will make global observations of the inner heliosphere with a constellation of four small satellites (*PUNCH - NASA Science*, 2023). The primary PUNCH measurement is brightness and polarization state of light scattered by electrons entrained in solar wind features (Killough et al., n.d.). This data is then processed on the ground to generate calibrated, fully mosaicked polarized images. Four satellites are necessary because PUNCH views the entire inner solar system continuously, and Earth blocks approximately half the view from any one location in orbit (*PUNCH: The Mission*, n.d.). The spacecraft are positioned in a polar low Earth orbit, spread out near the terminator line, giving them a clear view in all directions around the Sun. By conducting synchronized observations across four satellites, PUNCH will be able to measure solar wind globally around the star in near real time. On April 14th, 2025, PUNCH captured its first image of the Sun. On August 7th, 2025, all four spacecraft successfully maneuvered into their final science orbits (*NASA's PUNCH Mission Reaches Science Orbit, Releases Data - NASA Science*, 2025).

The combination of wide field of view, high imaging cadence, and continuous observation makes PUNCH a standout mission in the state-of-the-art. However, its multi-spacecraft capabilities have yet been demonstrated in orbit.

Several commercial imaging constellations, including **Planet Labs** and **Pleiades Neo**, further demonstrate distributed observation capabilities. While typically not coordinated as swarms, these systems contribute relevant technologies such as high-resolution optical sensing, automated tasking, and persistent monitoring from orbit (*Airbus to Reshape Earth Observation Market with Its Pleiades Neo Constellation*, n.d.; *Tasking API | Planet Documentation*, 2025; Boshuizen et al., 2014).

Perception

Perception is a multi-agent sensing capability that enables agents to interpret their surroundings using raw or processed sensor measurements. It encompasses detection and recognition, where agents identify and classify objects or feature of interest, as well as mapping the environment. This section reviews the state of the art in perception for multi-agent systems, focusing on mission technologies that have been demonstrated in flight. When relevant, flown single-agent or planned multi-agent missions are discussed to highlight areas where future swarm missions could offer meaningful advancements. The overall Sensing Capability of each mission is qualitatively evaluated using the metrics identified earlier.

Among the most advanced perception systems demonstrated in a flown multi-agent space system is the **Starling Formation-Flying Optical Experiment (StarFOX)**, one of the payloads of the NASA Starling mission, which consists of four CubeSats in LEO launched in 2023. StarFOX applies the angles-only Absolute and Relative Trajectory Measurement System (ARTMS) to perform onboard target detection, recognition, and tracking using star tracker cameras. The image processing algorithm (IMP) in ARTMS first computes a set of internal bearing angles for each target in a captured image. Resident space objects are identified in images by matching them to known objects from a space catalog. The computed bearing angles are then used to track known targets and detect new targets. This represents a high degree of onboard perceptual autonomy in a disturbed flight environment. StarFOX ARTMS captures approximately 50-100 image per orbit and has been shown in flight to track as many as two targets simultaneously. However, the star tracker images are noted to have poor signal-to-noise (SNR) ratio, thus targets at distances greater than 100 km could not be detected (Kruger et al., 2023).

The **Cyclone Global Navigation Satellite System (CYGNSS)** mission, launched in 2016, consists of a constellation of eight small satellites designed to improve hurricane forecasting by measuring ocean surface wind speeds (Carreno-Luengo et al., 2021). Each satellite collects reflected GPS signals to generate Delay Doppler Maps (DDMs), which capture the spatial distribution of ocean surface roughness and enable the estimation of near-surface wind speeds. The raw DDMs are produced onboard each CYGNSS satellite and downlinked to the ground where they are further processed to produce Level 1 (L1) and Level 2 (L2) products. The mission has a revisit frequency of approximately 7 hours and produces four DDMs per second per satellite in its primary operating mode. This multi-satellite sensing capability provides dense spatial and temporal coverage across the tropics that surpass what an individual spacecraft can accomplish. Thus, CYGNSS represents a state-of-the-art demonstration of multi-agent environmental mapping, delivering observation products critical for cyclone and ocean surface monitoring.

The **RADARSAT Constellation Mission (RCM)**, operated by the Canadian Space Agency, is a three-satellite constellation launched in 2019 for real-time maritime surveillance, ecosystem monitoring, and disaster management (Cote et al., 2021; Doyon et al., 2018). Each satellite is equipped with a Synthetic Aperture Radar (SAR) and Automatic Identification System (AIS) payload. SAR delivers high-resolution ($\geq$ 1 meter) radar imaging for detailed monitoring of ship movements and maritime activities. AIS complements the SAR data by receiving identification messages broadcast by the vessels, which enables ship detection and recognition. The three-satellite architecture offers frequent revisit and wide-area coverage, daily access to 95% of any location on Earth (*RADARSAT Constellation - eoPortal*, n.d.). Data is downlinked and processed on the ground to generate SAR image products. Together, these capabilities demonstrate a state-of-the-art multi-agent perception system for maritime observation and ship detection.

The **Gravier Recovery and Climate Experiment Follow-On (GRACE-FO)** and **Gravity Recovery and Interior Laboratory (GRAIL)** missions are landmark demonstrations of multi-agent gravitational field mapping of Earth and the Moon, respectively. Both missions utilized a pair of spacecraft flying in tandem, precisely measuring the distance changes between them as they respond to variations in the gravitational field of the planetary body they orbit. GRACE-FO, launched in 2018 as the successor to the original GRACE mission, uses inter-satellite ranging via a microwave instrument to detect minute shifts in the distance between the two Earth-orbiting satellites, which result from gravitational anomalies caused by changes in mass distribution on the ground (*GRACE-FO*, n.d.; Landerer et al., 2020). Similarly, the GRAIL mission, launched in 2011, employed the Lunar Gravity Ranging System (LGRS) to acquire high-precision range-rate measurements between two spacecraft orbiting the Moon. These measurements enabled the construction of the highest-resolution gravity map of the Moon to date, revealing insights into its crustal structure and thermal history (Edwards-Stewart, n.d.; Zuber et al., 2013). In both missions, the measurements between agents enabled them to generate high-resolution maps of gravitational fields that would be impossible to capture with a single-agent platform.

Electrojet Zeeman Imaging Explorer (EZIE) is a NASA heliophysics mission, launched in 2025, that will study the Sun and space weather near Earth. It consists of a constellation of three CubeSats to observe auroral electrojets – strong electric currents flowing through the polar regions of Earth – with the goal of discerning their spatial structure. EZIE temporally and spatially observes the electrojets by studying the emissions from oxygen molecules below them. In the presence of a magnetic field, such as those created by the electrojets, the emission lines split in a phenomenon called the Zeeman effect. This splitting reveals the strength and direction of the magnetic field. EZIE uses its Microwave Electrojet Magnetogram (MEM) instrument on each satellite to detect Zeeman splitting of the oxygen emission line. So far, one of the three satellites has captured “first light” and successfully recorded Zeeman splitting using the MEM instrument (Gjerloev et al., 2024; Misra et al., n.d.; *NASA’s EZIE Launches on Mission to Study Earth’s Electrojets - NASA Science*, 2025). The raw spectral data is downlinked and processed on the ground to deliver calibrated magnetic field maps and electrojet structure. Given its recent launch date, NASA has yet to publish experimental results from the flown mission, and its multi-satellite capabilities have yet been demonstrated.

Laser Interferometer Space Antenna (LISA) is a collaborative NASA-ESA mission scheduled to launch in the mid-2030s. It will consist of three identical spacecraft arranged in an equilateral triangle with 2.5 million kilometers of separation, trailing Earth in its orbit around the Sun by approximately 50 million kilometers. This formation enables LISA to detect gravitational waves with an interferometric measurement of differential optical pathlength modulation along the three sides of the triangular configuration. The laser interferometer will continuously measure minute changes in the distances between each spacecraft at frequencies from 0.1 millihertz to 1 hertz, with an accuracy of approximately 10 picometers, which is equivalent to detecting the diameter of a proton across millions of kilometers (Amaro-Seoane et al., 2017; *LISA – Laser Interferometer Space Antenna – NASA Home Page*, n.d.). This accuracy enables LISA to detect gravitational wave events, such as supermassive black holes mergers, across all ages of the universe, from cosmic dawn to the present. Once deployed, LISA will become the first gravitational wave detector in space (*LISA Factsheet*, n.d.). By leveraging a coordinated constellation of spacecraft with highly sensitive interferometers, LISA exemplifies multi-agent perception in space.

Though not multi-agent, single-spacecraft missions such as the NASA **Cislunar Autonomous Positioning System Technology Operations and Navigation Experiment (CAPSTONE)** and the ESA **OPS-SAT**, launched in 2022 and 2019 respectively, have demonstrated advanced onboard detection and recognition capabilities using machine learning. Specifically, the Sigma Zero system on CAPSTONE successfully performed anomaly detection and classification of navigation estimates in lunar orbit using a neural network (NN) (Gardner, Cheetham, Bolliger, et al., n.d.; Gardner, Cheetham, Parker, et al., n.d.), while OPS-SAT successfully demonstrated onboard image classification using a convolutional neural network (CNN), image clustering with K-Means, and fault detection, isolation, and recovery (FDIR) using a support vector machine (SVM), all in Earth orbit. The FDIR model, in particular, trained the AI model onboard the spacecraft (Labrèche et al., 2022). In all other cases, the machine learning models were trained on the ground and uplinked to the spacecraft for onboard inference. These efforts illustrate autonomous perception capabilities that could potentially translate to the next generation of multi-agent space missions.

Awareness

Awareness in multi-agent sensing encompasses estimating both the internal and external states of an agent from inputs generated by distributed sensors. Key capabilities include localization, such as attitude estimation and orbit determination. This section reviews the state of the art in awareness for multi-agent systems, focusing on mission technologies demonstrated in flight. Where relevant, flown single-agent or planned multi-agent systems are also discussed to highlight capability gaps and potential directions for future swarm missions. The overall Sensing Capability of each mission is qualitatively evaluated using the metrics identified earlier.

Precise estimation of position and orientation is essential for awareness. The **StarFOX ARTMS** payload on the NASA Starling mission, mentioned previously, represents the most advanced demonstration of this capability. StarFOX uses angles-only measurements from star tracker cameras to perform onboard absolute and relative orbit determination of nearby resident space objects. Using only bearing angles to optical targets – without GPS – StarFOX achieved a relative position uncertainly of 1.3% of the target range for one observer and 0.6% with multiple observers (Kruger et al., 2024). This represents the first in-orbit demonstration of autonomous angles-only absolute and relative navigation for a satellite swarm and establishes a strong foundation for future multi-agent missions in GPS-denied environments.

While not a multi-agent system, the **CAPSTONE** mission, mentioned previously, provides a compelling demonstration of autonomous localization in deep space. Launched in 2022, CAPSTONE includes the Cislunar Autonomous Positioning System (CAPS), which enables the spacecraft to estimate its orbital state using crosslink ranging to the NASA Lunar Reconnaissance Orbiter (LRO) satellite. During a CAPS demonstration in May 2023, CAPSTONE transmitted a ranging signal to LRO and received a reply, allowing it to compute its position and velocity autonomously without relying on GPS. This onboard processing achieves localization uncertainties on the order of kilometers in position and centimeters per second in velocity – comparable to traditional ground-based methods – and enables station-keeping maneuvers in the Near-Rectilinear Halo Orbit (Gardner, Cheetham, Bolliger, et al., n.d.; Gardner, Cheetham, Parker, et al., n.d.). By successfully implementing real-time, peer-to-peer navigation in cislunar space, CAPSTONE marks a significant step forward in spacecraft autonomy and serves as a foundational example of how localization can be distributed across agents in GPS-denied environments.

Summary

This chapter discussed the current state of the art in Multi-Agent Perception. Key capabilities in data collection, perception, and awareness were evaluated across missions, using the identified the Sensing Capability Metrics – accuracy, speed, resolution, and coverage.

In data collection, the state-of-the-art missions include:

- **COSMIC-2**, which produces high-volume radio occultation measurements for ionospheric and atmospheric profiling. This mission demonstrates low-latency, high-precision environmental data collection across a constellation.

In perception, the state-of-the-art missions include:

- **StarFOX**, which showcases advanced onboard detection, recognition, and tracking of nearby targets using star tracker imagery and derived bearing angle measurements. This marks a key milestone in autonomous, multi-agent perception.
- **GRACE-FO** and **GRAIL**, which demonstrated multi-agent gravity field mapping of Earth and the Moon respectively, through inter-satellite ranging. These missions enabled planetary-scale awareness unachievable by a single agent.
- **CYGNSS**, which uses GPS-reflected signals to deliver dense wind measurement maps for hurricane monitoring, demonstrating real-time, multi-agent environmental mapping from orbit.

In awareness, the state-of-the-art missions include:

- **StarFOX**, again, for its autonomous relative and absolute orbit determination using angles-only optical navigation – a first of its kind in orbit.

Despite these advances, critical gaps remain in Multi-Agent Sensing. A primary limitation is the level of coordination across agents. Most of the missions discussed here operate as a constellation, with each agent either behaving or operated independently of others in the multi-agent system. Satellite constellations provide increased spatial and temporal coverage, but they lack the collaborative interaction capabilities needed to share information. This limits their ability to meaningfully sense themselves and their surrounding environment, and further, hinders these missions from fully leveraging the advantages of a multi-agent architecture. Building upon this limitation, another significant challenge lies in achieving robust sensor fusion across agents to improve perception and awareness capabilities. There currently exists no mission that has successfully demonstrated multi-agent sensor fusion in flight, despite its potential for increased accuracy, broader coverage, and improved resilience.

Furthermore, while there have been multiple missions to achieve distributed localization and mapping around Earth, these capabilities have yet to be extended to other planetary bodies. Such a development is crucial for future deep-space exploration and autonomous operations. Additionally, the full potential of onboard detection and recognition capabilities using machine learning across multiple agents has yet to be fully realized, offering avenues for enhanced autonomy and real-time decision making. Lastly, there remains a gap in target tracking with multiple observers. While StarFOX shows promise, it has so far only demonstrated target tracking with at most two observers, and importantly, relies on a coarse initialization of the target orbit to achieve convergence. These limitations highlight a need for scalable multi-agent systems that can leverage observations from multiple agents to improve accuracy and convergence speed. Addressing these gaps will be vital to enabling multi-agent systems that can effectively sense their environment in Earth orbit and beyond.

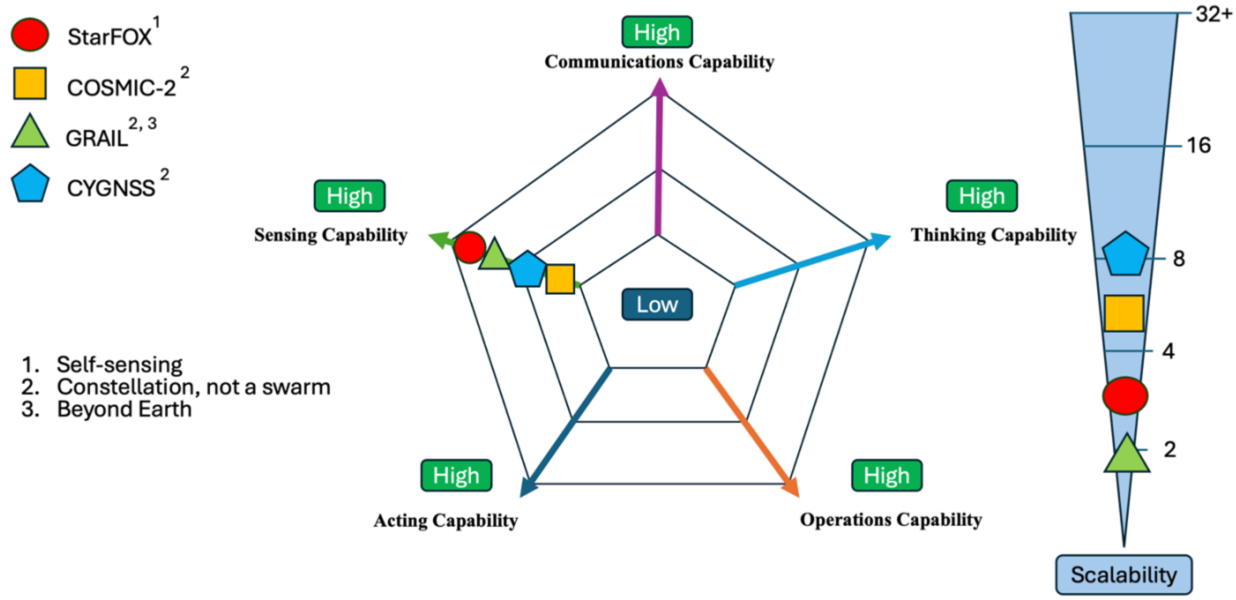

Figure 4: A radar chart of the five domains in multi-agent systems, with the state-of-the-art multi-agent sensing missions plotted. Moving farther out radially corresponds to a more capable system.

Multi-Agent Thinking

Introduction

Multi-Agent Thinking refers to the collective capabilities that enable a multi-agent system to select coordinated actions or joint plans that achieve system-level goals. These capabilities form the foundation for distributed planning, learning, prediction, and fault management, allowing multiple spacecraft to operate as an intelligent and adaptive team. Rather than relying on centralized control or pre-determined behaviors, multi-agent thinking systems negotiate goals, allocate tasks, forecast the intent of other agents, and adapt plans based on evolving mission contexts. In space-based applications, this includes joint mission planning, dynamic scheduling, and adaptive fault management.

Key Capabilities

- Multi-Agent Planning and Scheduling:
  - Goal and task negotiation technologies support agreement on current and future activities, their priorities, and their disposition among elements or systems.
  - Mission planning and task allocation can be thought of as 'strategic,' while activity planning and scheduling, and path and motion planning, are 'tactical.'
  - Communications planning and scheduling is a distinguished capability, as it may be needed to perform other multi-agent mission functions.
- Joint Knowledge and Understanding:
  - Technologies that support collection, assembly, sharing, and interpretation of information and intent among elements to solve problems and plan actions/responses.
  - Includes multi-agent sensing, sense-making such as map making, uncertainty estimation and target prioritization.
- Behavior and Intent Prediction:
  - Technologies that forecast the actions of other elements or systems to support collaboration and interaction.
  - Includes inferring plans and actions of other agents by observing their actions.
- Multi-Agent Fault Management:
  - If one member of the multi-agent system has failed and can't diagnose its own failures, other members of the system may be able to do so, either by means of indirect information and data gathering (e.g., monitoring communications), or by performing directed observations (e.g., planned in-space inspection).

For a single agent, thinking capabilities include planning, learning, and fault management. **Mission planning and scheduling** technologies optimize tasks based on the goals, objectives, and activities needed to achieve a mission, subject to the situation, and constraints. Planning and scheduling capabilities also select and order tasks and activities, while managing system resources and satisfying constraints, in order to achieve mission goals. **Path planning** technologies generate or modify a path or trajectory to reach a desired target physical location or physical configuration, subject to system and environment constraints. **Fault diagnosis and prognosis** technologies detect and identify faults, predict future faults, and assess system capability as a consequence of faults that have occurred. **Learning** technologies adapt to changing environments and conditions without explicit re-programming an agent's own experiences.

Extending these capabilities to multiple agents introduces new dimensions to thinking capabilities. Groups of agents can think in a variety of ways. "Cooperative" agents have aligned goals, sometimes described as "shared reward signal" in multi-agent learning. "Coordinated" agents interact to plan and execute joint actions (i.e., their actions impact each other, so they don't have fully decoupled decision problems), but agents may pursue their own objectives. Multiple agents can (and must) share information to jointly define the problem to be solved, be it information collected about the environment or their own capabilities. In some circumstances, multiple agents can reduce thinking workload by "dividing and conquering" complex tasks amongst different agents. When deciding on courses of action, agents can take advantage of their different capabilities, including design (different sensors) or accumulated degradation or faults (loss of sensors). When learning, agents can learn using their own experiences, or the experiences of other similar agents. Finally, it is notable that agents may need to think about how to communicate in order to perform these other multi-agent mission functions.

The history of thinking systems in space starts with the Remote Agent Experiment (RA) in 1998, which featured both automated planning and fault management capabilities. The state of the art in autonomous thinking capability today is represented by the task, activity, and path planners onboard the Curiosity and Perseverance rovers, currently operating on Mars, and the automated planning and fault management technologies that are slated to fly on Gateway as part of the Vehicle Systems Manager (VSM). While machine learning classifiers have flown on both Perseverance and Curiosity to assist in targeting robotic sampling, as well as on Earth-Observing One to classify cloudy pixels and identify events such as volcanoes and floods, only recently have machine learning algorithms been executed in space. However, there are few examples of in-space multi-agent systems, and only one that is fully decentralized.

Thinking Capability Metrics

- Decision Distribution
- Temporal Responsiveness
- Planning Depth
- Optimality

To evaluate the state of the art of Multi-Agent Thinking, this report introduces a set of metrics which capture the performance of thinking capabilities in multi-agent systems. These metrics are qualitatively combined to assess the **Thinking Capability** of relevant missions and technologies. **Decision distribution** represents the spectrum of fully centralized to fully decentralized. For example, NASA's HelioSwarm mission employs a centralized approach, where a hub of spacecraft aggregates information and sends global decisions from ground to the other spacecraft. By contrast, NASA's Distributed Spacecraft Autonomy (DSA) system demonstrates a decentralized approach, with nodes independently collecting science data and reaching consensus on observation strategies through inter-satellite communication, without relying on a central controller. As these examples illustrate, decision distribution is impacted by both communication efficiently and data exchange. **Temporal responsiveness** reflects how quickly the system reacts to an event, and to an extent, is also impacted by communication efficiently and data exchange. **Complexity of decision** measures how large of a problem is solved. It is relevant to multi-agent thinking capabilities such as fault management and learning. More complex decisions require solving larger problems with more constraints, which in turn, requires more resources to devote to thinking. Similarly, more complex behavior to learn requires more data and more sophisticated learning algorithms, requiring devoting more resources to learning. An example of a complex decision is a swarm collaboratively diagnosing a fault in one spacecraft, redistributing sensing tasks across heathy agents, and replanning path trajectories to preserve mission objectives. **Optimality** refers to finding the best possible solution to a problem. Optimizing is hard even for centralized thinking systems, and can be even more challenging in a fully decentralized system. For our purposes, we assess then run-time plan quality trade-off.

Together, these metrics provide a framework for evaluating Multi-Agent Thinking capabilities. They offer insight into how effectively a multi-agent system can distribute decision-making, respond to events, plan over time, manage faults, and adapt through learning. They also highlight trade-offs between performance, robustness, and complexity, helping to identify strengths and limitations in state-of-the-art technologies. This section primarily focuses on flown multi-agent missions that have demonstrated the identified action capabilities. However, select examples from planned multi-agent missions are included when they illustrate technologies relevant to multi-agent thinking systems. In cases where no multi-agent system has flown yet for a given capability, the gaps are noted to highlight areas for future advancement.

State of the Art – Complete Swarm Systems

The **Distributed Spacecraft Autonomy (DSA)** system demonstrated collaborative resource allocation for multi-point science data collection as a software payload on NASA's Starling mission, consisting of four small spacecraft. Central to DSA's operations is an autonomous GPS channel selection process used to optimize channel selection across the spacecraft swarm; the goal is to demonstrate fast, autonomous reactions to uncertain events, optimizing the collection of relevant scientific data. DSA demonstrated the ability to capture ionospheric phenomena such as the Equatorial Ionization Anomaly and Polar Patches by utilizing relative Total Electron Content (TEC) measurements as sensor inputs. The DSA flight software architecture consists of three apps written in NASA's Core Flight System framework. The Comm App enables message routing over Starling's Ad-Hoc Network. The TEC App processes GPS receiver data from the spacecraft bus and produces inputs to the Autonomy App based on the spacecraft geometry and ionospheric features extracted from the TEC signal. The Autonomy App fuses the rewards from the local TEC App with the rewards from the other spacecraft, and uses a Mixed Integer Linear Program (MILP) solver to autonomously select the best set of GPS channels to monitor. All apps run on all spacecraft in a completely decentralized manner, enabled by a 'distributed consensus' approach, ensuring all spacecraft share a 'common operating picture' when selecting the next course of action. The DSA system's autonomous reconfiguration ability is demonstrated, emphasizing its adaptability to natural phenomena without significant or direct tasking. The experiment's success demonstrates the potential for autonomous systems in future space missions, enabling spacecraft to operate independently and efficiently. (It can be argued that DSA's capabilities cuts across Thinking and Acting, but the only 'actions' DSA takes are to extract TEC data from the different GPS satellites and communicate the needed information over Starling's Ad-Hoc network, so we have chosen to describe DSA under Thinking.)

The **Reconfiguration and Orbit Maintenance Experiments Onboard (ROMEO)** experiment, also demonstrated onboard the Starling spacecraft, is 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. CFA is a modular system that consists of the navigation service, which determines each satellite's location; 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 screened for conjunctions on the ground, and a final Go/No Go command was given. 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. ROMEO was demonstrated onboard Starling, but was only partly successful.

The **Cooperative Autonomous Distributed Robotic Exploration (CADRE)** mission is a lunar technology demonstration mission of multi-agent autonomy on a team of three rovers and a base station. The goal of CADRE is to demonstrate how a team of autonomous rovers, receiving only high-level tasks from Earth, can autonomously explore a region of the Lunar surface, as well as perform a distributed measurement in coordination with a multi-static ground-penetrating radar. CADRE uses a centralized planner on a designated lead rover to partition unexplored areas to individual rover agents, which in turn generate their own path plans independently. In order to perform its duties, the leader needs to collate information from all agents. In particular, the strategic planner requires knowledge of which agents are able to participate in multi-agent autonomy (i.e., which agents have been deemed operational by ground operators and on-board fault protection), and the power and temperature of all agents; in addition, the team planners for exploration and distributed measurements require knowledge of the agents' location and, critically, of a global, up-to-date map of the environment. A distributed database is responsible for collecting this data from every agent and synchronizing it to the leader and designated survivor. The distributed database interfaces with power and temperature estimators provided by flight software; with fault protection; and with surface mobility, which provides each rover's location and periodically output local maps centered around the robot's location. While CADRE is not a completely decentralized system, leader and survivor election are performed in a decentralized manner, and individual agents pose and solve their path planning problems in a decentralized manner as well. CADRE is slated to fly in 2026.

NODES was a pair of 1.5 U CubeSats developed by the NASA Ames Research Center under the Small Spacecraft Technology Program (SSTP) within NASA Space Technology Mission Directorate (STMD). The NODES spacecraft were launched to the ISS in December 2015 and deployed from the ISS in early May of 2016. NODES was designed to expand the utility of small spacecraft networks and to explore issues related to the command and control of swarms of multiple spacecraft making synchronized, multipoint scientific measurements. The NODES spacecraft autonomously self-organized, selecting which spacecraft will lead the formation, collect data and pass those data to the ground, all based on the states of the spacecraft. It also demonstrated the planning and commanding of spacecraft in a swarm, from the ground and through the network. Finally, NODES performed synchronized, multi-point science measurements of the Earth's charged particle environment. Similar to CADRE, NODES is a 'centralized' system inasmuch as a single spacecraft is always 'captain,' and thus a single agent plans for the remaining agent. However, NODES represents one of the earliest efforts to build a complete multi-agent system with explicit in-space coordination and communication between elements.

NASA's Jet Propulsion Laboratory (JPL) is leading the **FAME** project, an in-space flight demonstration leveraging existing commercial assets from satellite providers such as Open Cosmos, Loft, Planet, Aptos, and D-Orbit to demonstrate NASA's New Observing Strategy (NOS) (Chien, 2025). This effort includes other collaborators such as Ubotica, Hyspace and others. The demonstration will show the value added in emerging in-space capabilities including: onboard analysis of acquired data, leveraging available onboard edge compute, onboard rapid self-response, leveraging the knowledge generated from the onboard analysis to direct/reconfigure sensors on the same platform, rapid notification of the ground or other assets, leveraging the knowledge from onboard analysis, and 24/7 satellite communications link; rapid cross-tasking of other assets, wherein insights derived on one spacecraft are used to drive tasking of another spacecraft; and multi-agent systems technology to orchestrate user-defined workflows across both flight and ground and also across multiple spacecraft operated by different providers. Given the scale and heterogeneity of platforms, this effort is poised to become the state of the art for multi-agent thinking capabilities in Low Earth Orbit.

Summary

The two most mature in-space multi-agent thinking capabilities, DSA-Starling and CADRE, employ very different approaches. DSA-Starling uses a 'distributed consensus' approach that replicates the 'thinking' computation, task negotiation, across all agents in a fully distributed manner. CADRE, by contrast, uses a centralized approach to goal and task negotiation that allows fully decentralized (independent) path and motion planning. However, leader and survivor election and global database distribution require more distributed coordination.

Notably, no missions to date have demonstrated distributed or decentralized fault management or machine learning. Also notably, DSA-Starling's distributed consensus approach does not exploit divide-and-conquer strategies that could reduce run-time, but these approaches may be complex and are communication-intensive.

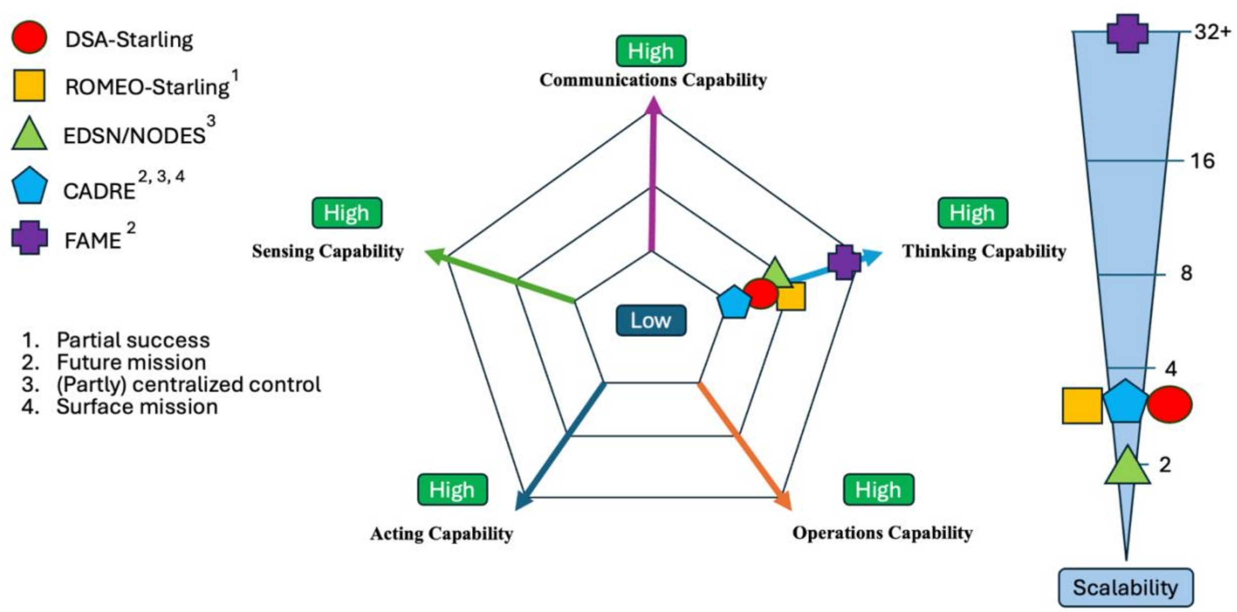

Figure 5: A radar chart of the five domains in multi-agent systems, with the state-of-the-art multi-agent thinking missions plotted. Moving farther out radially corresponds to a more capable system.

Multi-Agent Action

Introduction

Multi-Agent Action refers to the ability of a system of agents to achieve a goal through coordinated behaviors. This domain is fundamental for tasks requiring collaborative movement and adaptive response. Rather than individual agents acting in isolation, multi-agent action systems integrate the movements and decisions of multiple agents to achieve goals that would be impossible or inefficient for a single agent. In space-based swarms, specifically, agents often need to maintain formations, execute trajectories, and adapt to faults or changes in the mission environment. Effective action requires real-time control across agents, robust communication or shared state estimates, and mechanisms for autonomous response to internal or external disturbances. The goal of this domain is to ensure that multi-agent systems can move and act as a unit.

Key Capabilities

- Formation Flying
  - Motion Control
  - Trajectory Following
  - Provably Safe Trajectories
- Fault Recovery and Response

Key action capabilities in this domain enable agents to translate shared goals into physical execution. At the core of this is **formation flying**, which involves multiple agents coordinating their relative positions and velocities to maintain desired spatial configurations. This includes **motion control**, which enables individual agents to adjust their thrust and attitude to achieve desired positions and velocities. Motion control is the fundamental building block for all coordinated movement. **Trajectory following** builds upon this by enabling agents to follow coordinate paths through space, ensuring the entire formation adheres to a desired spatial and temporal progression. **Provably safe trajectories** ensure that the planned paths satisfy safety constraints and avoid collisions. This formal verification provides provable guarantees on the safety of planned trajectories, even in the presence of uncertainty. **Fault recovery and response** ensures the resilience and continued operation of the multi-agent system in the face of unexpected events or failures. Separate from fault detection, these capabilities encompass the ability of agents to autonomously reconfigure the system or redistribute tasks to maintain mission objectives. This is vital for maintaining system integrity and mission success in dynamic and unpredictable environments.

Action Capability Metrics

- Level of Coordination
- Execution Accuracy / Precision
- Responsiveness
- Resource Use

To evaluate the state of the art of Multi-Agent Action, this report introduces a set of metrics that capture the performance of action capabilities in multi-agent systems. These metrics are qualitatively combined to assess the Action Capability of relevant missions and technologies. **Level of coordination** assesses the degree to which agents in a system work together to achieve a common goal. It ranges from loosely coupled systems where agents act largely independently, to tightly integrated systems where agents dynamically adjust their behaviors in response to others. Higher levels of coordination often imply more sophisticated inter-agent communication and decision-making. **Execution accuracy and precision** reflects how closely the actual movements and behaviors of the agents match the commanded or desired actions. High accuracy ensures that formations are maintained, trajectories are precisely followed, and collaborative tasks are executed without significant deviation. **Responsiveness** encompasses the speed with which the multi-agent system can react to new commands, environment changes, or unforeseen events. High responsiveness allows the swarm to quickly adapt its actions, reconfigure its formation, or initiate recovery procedures in real-time, which is critical for operating in dynamic and unpredictable environments. **Resource use** evaluates how efficiently agents utilize available resources, such as propellant, energy, or computing power. Optimizing resource use ensures that multi-agent systems can sustain operations for longer durations, handle more complex tasks, and maintain performance without comprising mission objectives.

State of the Art – Complete Swarm Systems

This section reviews the current state of the art in action capabilities for multi-agent systems – specifically formation flying and fault recovery and response – with a focus on flown missions that involve more than one spacecraft. Landmark missions such as PROBA-3 have achieved unprecedented levels of precision in formation flying in orbit. Looking ahead, the Cooperative Autonomous Distributed Robotic Exploration (CADRE) mission, though currently in its planning and ground-testing phases, represents a pivotal step towards multi-agent action for lunar surface exploration. While some multi-agent architectures inherently provide a level of fault tolerance, inasmuch as the loss of one agent may not end a mission, multi-agent fault recovery and response in the space environment remains an active area of research.

This report focuses exclusively on multi-agent capabilities that have been successfully demonstrated in actual flight by space systems, encompassing both satellites and planetary rovers. The broader field of multi-agent robotics includes extensive research and real-world demonstrations with terrestrial Unmanned Aerial Vehicles (UAVs), such as complex formation flight, object manipulation, obstacle avoidance, and fault tolerance demonstrations by programs like the Defense Advanced Research Projects Agency (DARPA) OFFensive Swarm-Enabled Tactics (OFFSET) (Abro et al., 2024; *OFFensive Swarm-Enabled Tactics*, n.d.). While these systems illustrate multi-agent action capabilities, they are not in the scope of spacecraft flight demonstrations.

Formation Flying

Formation flying is a multi-agent action capability that involves two or more spacecraft maintaining precise positions and orientations. This coordinated control allows the individual spacecraft to function collectively as a single, distributed system and achieve complex mission objectives beyond what is possible for a single agent. The missions detailed in this report represent the state of the art in formation flying for multi-agent systems. A focus is on technologies that have been demonstrated in flight, but flown single-agent or planned multi-agent missions are discussed when relevant to highlight areas where future multi-agent missions could offer meaningful advancements. The overall Action Capability for the capabilities of each mission are qualitatively evaluated using the metrics identified earlier.

The ESA **Project for Onboard Autonomy-3 (PROBA-3)** mission stands as a landmark in multi-agent formation flying in orbit. Launched in 2024, PROBA-3 comprises of two spacecraft, the Occulter and the Coronagraph, and is designed to create an artificial solar eclipse to study the faint corona of the Sun. The system successfully maintained a formation at a distance of 150 meters with millimeter precision for several hours. This unprecedented level of precision was critical for the scientific objectives of the mission, allowing the Coronagraph to operate within the shadow cast by the Occulter. The autonomous maintenance of this precise positioning relies on a suite of sensors. GPS information was used to determine the location of the two satellites in space, while cold gas thrusters were commanded to eject minute amounts of propellant for fine adjustments to maintain formation. An intersatellite link facilitated the exchange of vital positioning information between agents, enabling the Occulter to send instructions to the Coronograph, and allowing for autonomous adjustments if misalignment is detected (Peñin et al., n.d.; *PROBA-3 (Project for On-Board Autonomy-3) - eoPortal*, n.d.; TIRAPLEGUI et al., 2019). The level of precision PROBA-3 achieved places it at the top of the current state of the art for multi-agent formation flying.

The NASA **Magnetospheric Multiscale (MMS)** mission has also successfully demonstrated multi-agent formation flying. Launched in 2016, MMS comprises four identical spin-stabilized spacecraft to study the magnetosphere of the Earth. The spacecraft fly in tetrahedron formations of varying dimensions to perform magnetospheric science measurements, with separations from tens of kilometers to a few hundred kilometers. MMS was the first mission to conduct a definitive study of magnetic reconnection in space, a phenomenon which drives space weather in the magnetosphere of Earth. To determine the orientation, dynamics, and impacts of magnetic reconnection diffusion regions, the spacecraft travel in a long, oval-shaped orbit around Earth, repeatedly passing through key regions where magnetic reconnection occurs. The formation maneuvers are executed by autonomously pulsing the thrusters to produce the targeted delta-v. MMS achieved millimeter-per-second velocity accuracy in magnitude and arc-minute directional precision (Chai et al., 2015; Williams et al., 2016). These capabilities put MMS at the forefront of multi-agent formation flying.

The University of Toronto's **CanX-4** and **CanX-5** mission were designed to demonstrate precise formation flying and inter-satellite communication at an ultra-low cost, capitalizing on the inherent advantages of nanosatellites. The mission launched in 2014 and comprised of a pair of identical nanosatellites, CanX-4 and CanX-5. The spacecraft used onboard propulsion, an S-band intersatellite communications link, and relative navigation using carrier-phase differential GPS techniques to perform a series of precise, controlled, autonomous formations from 1 km range down to 50 m separation. Three pieces of navigation and control software were used to full-fill the formation control requirements – the formation flying integrated onboard nanosatellite algorithm (FIONA), which computed formation keeping and reconfiguration control maneuvers and performed absolute state estimation, the relative navigation algorithm (RelNav), which was responsible for estimating the relative orbital state of the two spacecraft, and the onboard attitude system software (OASYS), which was responsible for re-orienting the spacecraft to commanded attitude targets and selecting the thruster(s) to be used for upcoming maneuvers. The CanX-4 and CanX-5 mission successfully demonstrated a series of autonomous formations with sub-meter control and centimeter-level relative position knowledge (Bonin et al., n.d.), pushing the boundary of what can be achieved with nanosatellites.

The Australian **M2** mission comprises of two 6U CubeSats owned and operated by the University of New South Wales (UNSW). Launched in March 2021, M2 aimed to demonstrate close proximity satellite formation flying in LEO and became the first mission to perform propellant-less formation flying through differential aerodynamic drag control (Brown et al., 2023; *M2 Mission | UNSW Canberra Space - UNSW Canberra*, n.d.). Traditionally, satellites alter their orbit through control of active onboard propulsion systems. However, differential aerodynamic drag enables the satellites to control their relative motion without the need of an active propulsion system (*Satellite Operations and Formation Flying | UNSW Canberra Space - UNSW Canberra*, n.d.). Theoretical concepts for the use of differential aerodynamic drag control for close proximity operations have been reported in the literature, but M2 marks the first on-orbit demonstration. Using an open "human in the loop" control approach, UNSW successfully reduced the along-track separation of M2-A and M2-B from 160 km to hold a steady along-track separation distance of 2 km for 3 days. Brown et al. have also proposed an autonomous closed-loop formation control system for close-proximity station keeping, to replace the time-intensive human-in-the-loop control approach currently used on M2. However, the proposed method has so far only been tested in simulation. Nonetheless, the successful demonstration of propellant-less close-proximity formation flying makes M2 a state-of-the-art mission in this domain.

Space Weather Atmospheric Reconfigurable Multiscale Experiment (SWARM-EX) is a National Science Foundation (NSF) sponsored CubeSat mission, scheduled to launch in 2026. The primary goal of SWARM-EX is to study the spatial and temporal variability of ion-neutral interactions in the equatorial Ionosphere-Thermosphere (I-T) region (*Project: Space Weather Atmospheric Reconfigurable Multiscale Experiment (SWARM-EX)*, n.d.). The mission consists of three 3U CubeSats operating in LEO that will take in-situ measurements of neutral and ion composition on timescales of less than an orbital period to study the persistence and correlation between different phenomena in the I-T region. The engineering objectives of SWARM-EX are focused on advancing the state of the art in spacecraft formation flying. Specifically, the mission will operate as a swarm and autonomously reconfigure its formation with minimal propellant consumption by utilizing a novel hybrid propulsive / differential drag control methodology (Agarwal et al., n.d.). The primary orbit reconfigurations require that the spacecraft maintain mean along-track separations from < 10 km to 1,400 km (Lowe et al., 2024). The proposed system will require minimal input from ground control, provide passive safety at close inter-spacecraft separations, and efficiently achieve large swarm reconfigurations. Once successfully demonstrated, SWARM-EX will be the first mission to perform fully autonomous formation flying with more than two spacecraft. Thus, it stands at the forefront of multi-agent action capabilities.

The NASA **Cooperative Autonomous Distributed Robotic Exploration (CADRE)** mission represents the cutting edge in multi-robot autonomy for planetary exploration. While its flight demonstration is slated for late 2025, its extensive ground testing signify a critical step in advancing multi-agent action capabilities in space. CADRE aims to send a trio of suitcase-sized rovers to the Moon, specifically designed to demonstrate that multiple robots can cooperate and explore together autonomously. A key objective is to follow path trajectories and collect distributed ground-penetrating radar (GPR) measurements to map the lunar subsurface (De La Croix et al., 2024). The rovers are designed to be largely autonomous – making decisions and acting without direct human intervention from Earth and communicating via mesh network radios with each other and a base station located on the lunar lander. In ground-test experiments, the rovers achieved an explore efficiency of 2.4 m^2/m (explored area / meters driven) and an explore rate of 0.08 m^2/s (explored area / duration) in a test area of 15 m x 22.5 m, with prototype rovers in an environment that is representative of the lunar surface (Nayak et al., n.d.). These results outperform both fully centralized and semi-centralized algorithms, and further, validate the performance of CADRE under mission-representative conditions. When successfully demonstrated in flight, CADRE will be a new benchmark in multi-agent trajectory following in space.

Fault Recovery and Response

The ability of multi-agent systems to recover from faults and respond adaptively is a critical advantage, enhancing mission resilience and continuity. While some multi-agent space missions have demonstrated ground-controlled reconfigurations in response to anomalies, the explicit demonstration of autonomous multi-agent fault recovery and response in flight remains an open problem. In most cases, multi-agent missions use the number of nodes as the primary fault response – if one node fails, the system simply continues operating with the remaining members. In such scenarios, there is no formal autonomous fault recovery procedure in place. This section reviews the state of the art in fault recovery and response for multi-agent systems, with a focus on planned mission technologies, as no multi-agent fault recovery system has yet been demonstrated in flight. The Action Capability for the capabilities of each mission are qualitatively evaluated using the metrics identified earlier.

The main mission of interest in multi-agent fault recovery is the **Virtual Super Optics Reconfigurable Swarm (VISORS)** mission – a distributed telescope consisting of two 6U CubeSats that will perform precision formation flying to detect and study energy release regions of the solar corona. The VISORS mission includes a formation level fault management system, which detects and responds to faults. The system monitors telemetry from all subsystems on both spacecraft and determines if any of them have crossed their nominal thresholds. The key to the diagnosis is that fact that the system has full state knowledge on the health of the formation. Once the fault has been diagnosed, the system responds to the fault via a combination of a mission mode switch, a subsystem power cycle or shut down, or a spacecraft role switch. Regardless of the fault diagnosis passed into the fault response module, the response is always chosen to be as conservative as possible. This keeps the logic simple and ensures that the spacecraft can make any required autonomous decisions to keep the formation safe, while also waiting for ground operators to handle more complex formation-level responses (Arunkumar et al., n.d.; Koenig et al., 2023).

Summary

This chapter discussed the current state of the art in Multi-Agent Action, highlighting key capabilities in both formation flying and fault response, and evaluating them with the identified Action Capability Metrics – level of coordination, execution accuracy, responsiveness, and resource use. In formation flying, the ESA **PROBA-3** mission represents that current state-of-the-art, having successfully demonstrated autonomous motion control among spacecraft. Despite not qualifying as a swarm, PROBA-3 achieved a separation distance between spacecraft of 150 meters, with millimeter precision, for several hours. To date, this is the closest formation ever flown in space. The **CADRE** mission, though not yet flown, represents the state-of-the-art in autonomous multi-agent trajectory following on planetary surfaces, showcasing high responsiveness and distributed coordination in ground testing. In contrast, autonomous fault recovery remains an emerging capability, with **VISORS** serving as a leading example of planned spacecraft-level fault response and recovery within a multi-agent architecture. Despite these advances, critical gaps remain in Multi-Agent Action. Specifically, no multi-agent mission has successfully demonstrated coordinated, autonomous fault response and recovery in flight, and multi-agent formation flying has only been demonstrated in LEO environments. These gaps indicate areas that the next generation of multi-agent swarm missions could fill, at the Moon, Mars, and beyond.

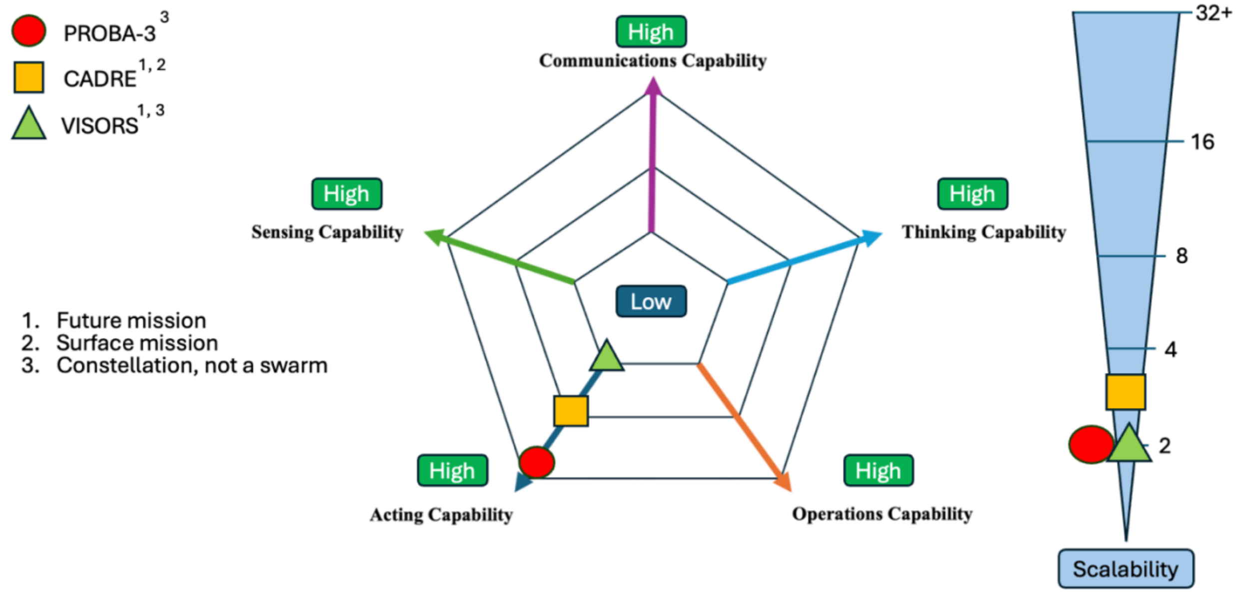

Figure 6: A radar chart of the five domains in multi-agent systems, with the state-of-the-art multi-agent action missions plotted. Moving farther out radially corresponds to a more capable system.

Multi-Agent Operations

Introduction

Multi-Agent Operations is concerned with how human operators interact with the swarm and to what extent they need to be involved in commanding the swarm. This domain was chosen as part of this study because operations is arguably one of the most important parts of a mission life cycle. Operations is where all of the design and engineering is put to the final test. It is where the science happens. Everything can go right, or everything can go wrong. Operations on a multi-spacecraft swarm mission can look very different from operations on a single spacecraft mission, especially as the swarm increases in number of agents. For this reason, it is critical to assess level of effort required to operate a swarm, how much agency is afforded to the swarm, and how operators are able trust the swarm is behaving correctly.

Key Capabilities

- Scalable Operations
- Behavior and Intent Prediction
- Operational Trust Building

Future swarms may scale to hundreds, perhaps thousands, of agents. At these sizes, it becomes infeasible for human operators to command each agent individually. Cost-effective operations require an operational paradigm that scales with the swarm. This capability relates to Taxonomy 10.3.5 Human and Machine Interactions. **Scalable operations** allow for intuitive and efficient interactions between the human operators and the swarm.

Scalable operations implies that some amount of work must be offloaded to the multi-agent swarm. If the agents are afforded a higher degree of autonomy, they must **behave predictably**. Operators of next generation multi-agent swarms must be able to predict what the swarm will do in several types of scenarios. The swarm must be designed such that it can be formally proven that it will behave as the operators expect and will not exhibit undesirable behaviors in off-nominal scenarios.

Allowing the swarm to have a higher degree of autonomy requires the operators to have a level of trust in the swarm. Trusting an autonomous system is a very careful balance. If operators have very little trust in the system, they will not use the swarm to its full capabilities, making operations extremely inefficient. On the other hand, if operators have too much trust in the swarm, the operators may allow the swarm to act beyond its capabilities, jeopardizing the mission. Operational trust building requires that the operator's level of trust matches the capabilities of the system. This idea is called **calibrated trust** (Lee & See, 2004). A next-generation autonomous swarm must provide operators with a degree of transparency that allows operators to build a level of trust that accurately matches the system's capabilities.

Operations Capability Metrics

- Workload Scale Rate
- Prediction Confidence
- Operational Transparency

Developing scalable operations is a complex task. The metric used to evaluate scalable operations is the **workload scale rate**, which represents the rate at which the operator workload increases as the number of agents increase. A simple next generation multi-agent swarm should have at minimum a multi-agent operations system that allows the number of agents in the swarm to increase without increasing the workload by the same factor. For example, increasing a swarm from ten agents to one hundred agents should not increase the workload by a factor of ten. In a more capable next generation multi-agent swarm, the workload would not increase at all with increasing the number of agents. The multi-agent operations system would be designed agnostic to the size of the swarm, allowing operators to command a swarm with any number of agents.

The metric used to evaluate behavior and intent prediction is the **level of confidence** operators have in their predictions of the system. A simple next generation multi-agent swarm will have undergone a minimal amount of testing to build up confidence in the predictability of the system. The testing will have covered the minimal scenarios for mission success and mitigation strategies in the event of off-nominal situations. A more capable next generation multi-agent swarm will have undergone much higher levels of testing. Testing occurs with finer granularity to accurately create a detailed picture of how the system will behave in numerous scenarios.

Operational transparency describes how an autonomous system is able to explain its confidence and lack of confidence in its abilities. A system that is **operationally transparent** can provide proof to the operator that it is functioning within its capabilities (Lyons, et al., 2017) (Lyons, et al., 2016). This proof allows the operator to put an amount of trust in the system, which leads to the operator's willingness to let the autonomy perform its duties. Additionally, an autonomous system that is operationally transparent will warn the user when it believes it may be operating outside its capabilities and explain its reasoning for this belief. This transparency allows the operator to carefully calibrate how much trust they are able to give the autonomous system, helping to ensure a careful balance of taking advantage of the autonomous system's capabilities without endangering the system by over-trusting it.

State of the Art – Complete Swarm Systems

NASA's **Starling 1.0** mission sits at the forefront of state-of-the-art multi-agent operations. The mission consists of four 6U CubeSats in LEO that demonstrated autonomous swarm operations. By employing an ad-hoc mesh network, Starling reduced the need for ground operators to command each spacecraft individually. In short, scaling from 1 to 4 satellites did not quadruple the operator effort – the workload scale rate remained low due to the on-board routing capabilities. Additionally, the Distributed Spacecraft Autonomy (DSA) software payload running on Starling supported the ability to autonomously plan operations, allowing the swarm to coordinate measurements and react to changes in the environment without continuous ground intervention. These capabilities make Starling 1.0 a standout mission in multi-agent operations.

Mars CubeSat One (MarCO) consisted of two identical 6U CubeSats that flew in tandem past Mars, alongside NASA's InSight mission, serving as a communications relay during the InSight lander's entry, descent, and landing on the Martian surface. The two CubeSats demonstrated a "carry-your-own-relay" functionality by receiving UHF telemetry from the InSight lander and forwarding it via high-gain antenna directly to Earth, with only the physical signal delay (~8 minutes) as latency. They demonstrated operational trust building by proving that small, low-cost spacecraft could reliably operate in the deep-space environment. Although MarCO did not include autonomy or coordination between agents, it was placed higher than on the operations capability scale because its demonstration environment – Mars approach – was significantly more demanding due to the high communication delay, radiation exposure, and limited power/propulsion compared to LEO missions.

PY4 (Stupl & Holliday, 2023) is a low-cost, rapid-turn-around technology demonstration based on the open-source PyCubed avionics framework. The mission consisted of 4 CubeSats in LEO that demonstrated over-the-air software/firmware updates and command forwarding. In principle, ground controllers only need to send a single uplink to one satellite, which then forwards that update or command to the others. This capability improves scalable operations by avoiding the need to individually command each spacecraft, keeping the workload scale rate low as the swarm grows. However, PY4 was rated lower on the operations capability scale because, while commands were forwarded automatically, the mission still required explicitly writing a command for each spacecraft. Still, PY4's demonstration of distributed command relays provides an important step toward resilient swarm autonomy, making it a notable mention in multi-agent operations.

Edison Demonstration of SmallSat Networks (EDSN) was a NASA technology demonstration mission that consisted of 8 satellites (Dalal, 2024). Its mission objective was to demonstrate networked operations for a swarm of satellites, with one spacecraft acting as a "Captain" to collect data from the others before sending it to a ground station. EDSN launched in November 2015 but was unfortunately lost due to a launch vehicle failure. Its follow-on, **the Network and Operation Demonstration Satellites (NODES)** mission, had the same technology demonstration goals but consisted of only two spacecraft. NODES was launched successfully in 2015 and deployed from the ISS in 2016. It reached full mission success and became the first mission to demonstrate spacecraft commanding via crosslink. Specifically, NODES demonstrated command routing from the ground to the "Captain" node to the "Lieutenant" node. NODES also utilized a scalable mission operations simulator, which allowed operators to evaluate how the system would perform under different swarm sizes and network configurations. Although the mission was limited to only two spacecraft, the demonstrated ability to route commands represents the first step toward scalable operations, making NODES a standout in the state of the art.

Summary

A next generation multi-agent swarm requires a multi-agent operations system that can allow the human operators to work efficiently and confidently with the swarm. The operations system must be scalable to prevent human operators from becoming overburdened. Additionally, the system must be predictable and provide operators with enough feedback to trust the system is operating correctly.

Figure 7 provides an overview for the capability of multi-agent operations. Overall missions in this domain remain low. MarCO is scored the lowest for multi-agent operations capabilities, because the agents are only capable of supporting operational trust building. They do not support functionality for scalable operations or behavior and intent prediction. Additionally, MarCO consists of only two agents. Py4 was scored slightly higher, as the command forwarding aids in scalable operations. Finally, Starling was placed highest on the radar gram because of its ability for cross commanding and DSA's autonomous planning capability.

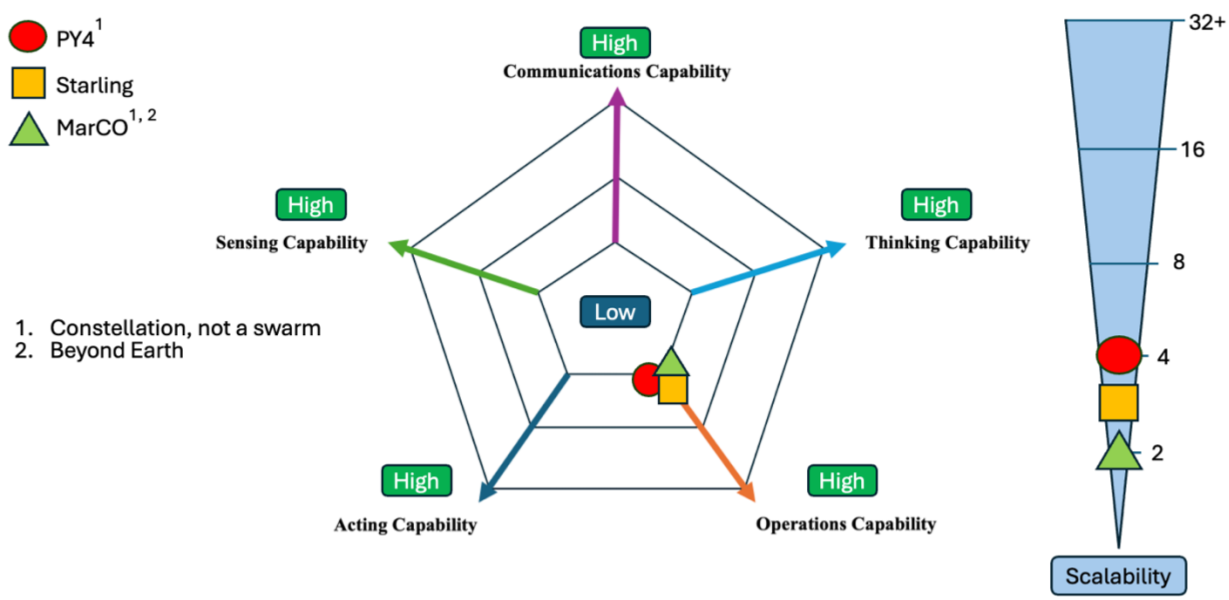

Figure 7: A radar chart of the five domains in multi-agent systems, with the state-of-the-art multi-agent operations missions plotted. Moving farther out radially corresponds to a more capable system.

Conclusion

The current state of the art in multi-agent swarm systems captures progress across the domains of communications, sensing, thinking, action, and operations. However, each domain also contains critical technology gaps that must be addressed to enable the next generation of multi-agent spacecraft swarms. This concluding section summarizes the top-performing missions in each domain, highlights existing limitations, and identifies key technology development needs that the next generation lunar swarm could adopt.

In **Multi-Agent Communications**, the SpaceX **Starlink** constellation represents the state of the art. With thousands of satellites in orbit, Starlink demonstrates the feasibility of large-scale mesh networking using laser communications. Despite not qualifying as a swarm, these capabilities dramatically reduce latency and increase resiliency, far surpassing any previous multi-agent space system. Thus, Starlink provides the blueprint for future swarm missions requiring high-throughput, low-latency connectivity. Gaps remain, however, in extending current mesh networking approaches, such as quality of service capabilities and interoperability. A next-generation multi-agent communications system must support interoperable and adaptive networking under challenging conditions, such as in lunar environments where communications relays are sparse and often constrained by limited bandwidth.

In **Multi-Agent Sensing**, several missions stand out among the state of the art. **StarFOX**, one of the core payloads of the NASA Starling mission, represents the first in-orbit demonstration of autonomous angles-only absolute and relative navigation for a satellite swarm. Its precise estimation of the position and orientation of targets, without the need of GPS, establishes a strong foundation for future multi-agent missions wanting to perform localization in deep space environments. Additionally, the NASA **GRAIL** mission, which consists of twin spacecraft in lunar orbit, performed gravity field mapping of the Moon through inter-satellite ranging. This mission produced the highest-resolution lunar gravity map to date, reveling insights into its crustal structure and thermal history. Although it does not qualify as a swarm, GRAIL demonstrated multi-agent sensing capabilities in the lunar environment, rather than around Earth, making it a notable standout. Together, StarFOX and GRAIL represent the current state of the art in multi-agent sensing capabilities. However, in each mission, the level of coordination across agents is limited and neither mission demonstrates on-orbit sensor fusion capabilities. Increasing collaborative interaction between agents will improve their ability to sense themselves and their surrounding environment. Likewise, incorporating sensor fusion will improve the situational awareness of the swarm and increase the accuracy of downstream tasks.

In **Multi-Agent Thinking**, the **DSA** software payload, flown on the NASA Starling mission, represents the current state of the art. Using a distributed consensus approach, DSA demonstrated real-time task allocation and adaptive sensing for ionospheric science, with spacecraft dynamically reconfiguring sensing tasks in response to changing conditions. This marks a milestone is distributed autonomy in space. However, critical gaps remain. No flown mission has demonstrated decentralized or distributed fault management, nor has a flown mission applied divide-and-conquer planning strategies to reduce run-time. Current approaches remain communication-intensive, and few examples exist of multi-agent learning executed in space. The next generation of multi-agent thinking systems must address these gaps to enable swarms to adapt more quickly, operate robustly under uncertainty, and make collective decisions with reduced reliance on ground intervention.

In **Multi-Agent Action**, the ESA **PROBA-3** and NASA **MMS** missions represent breakthroughs in coordinated autonomous formation flying in orbit. Although not technically swarms, these missions demonstrated a high degree of coordinated autonomy among multiple agents, making them standouts in the state of the art. Specifically, PROBA-3 demonstrated millimeter precision formation flying, maintaining a separation of 150 meters for several hours. An inter-satellite link facilitated the exchange of vital positioning information between agents, allowing for autonomous adjustments if misalignment was detected. Notably, PROBA-3 set the record for close-proximity formation flying in space. The MMS mission achieved millimeter-per-second velocity accuracy in magnitude and arc-minute directional precision in its formation flying demonstrations. Formation maneuvers were executed by autonomously pulsing the thrusters to produce the targeted delta-v. However, despite these achievements, critical gaps remain in multi-agent action systems. No multi-agent mission has successfully demonstrated distributed fault response and recovery in flight, and multi-agent formation flying has only been demonstrated in LEO environments. These gaps indicate areas that the next generation of multi-agent swarm missions could fill, at the Moon, Mars, and beyond.

Finally, in **Multi-Agent Operations**, the NASA **Starling** mission leads the state of the art, demonstrating scalable cross-commanding with multiple spacecraft. This represents the first step towards autonomous swarm operations that do not require a human in the loop. However, the field remains in its early stages – current operations architectures do not scale to large swarms without overwhelming human operators, nor do they provide thorough transparency into autonomous system decision-making. For these reasons, predictability and operator trust remain key challenges. Technology development must include scalable operations frameworks for human-swarm interaction and interfaces that provides transparency, explainability, and trust in autonomous decisions.

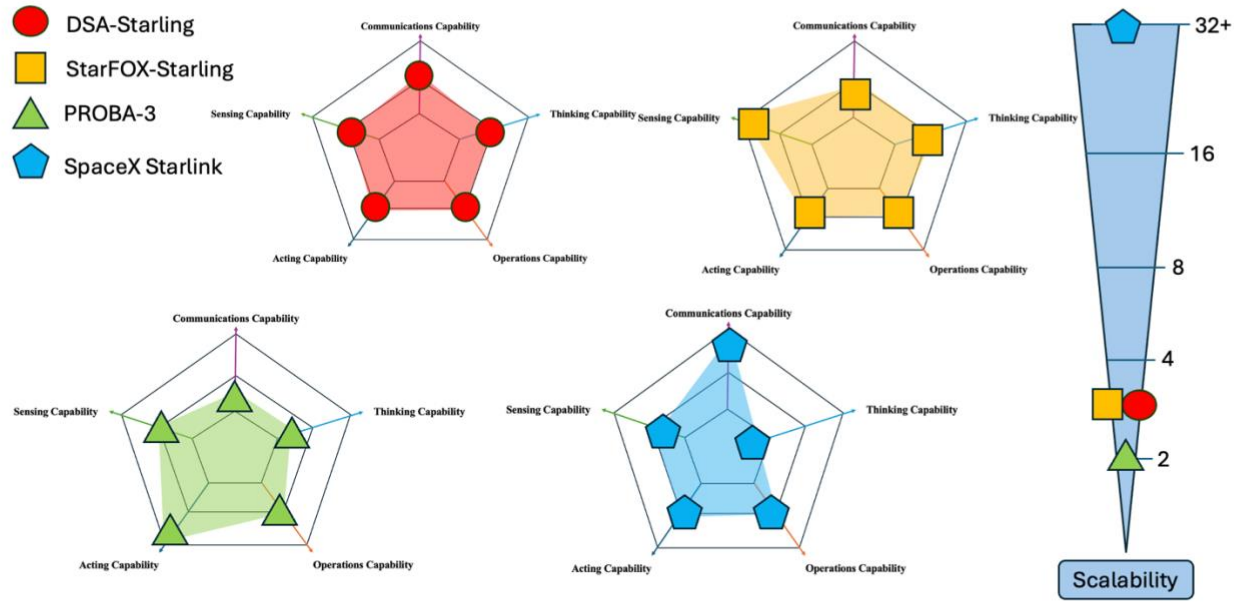

Figure 8: The top-performing multi-agent missions evaluated across all domains. The most capable missions will cover the most area on the radar chart.

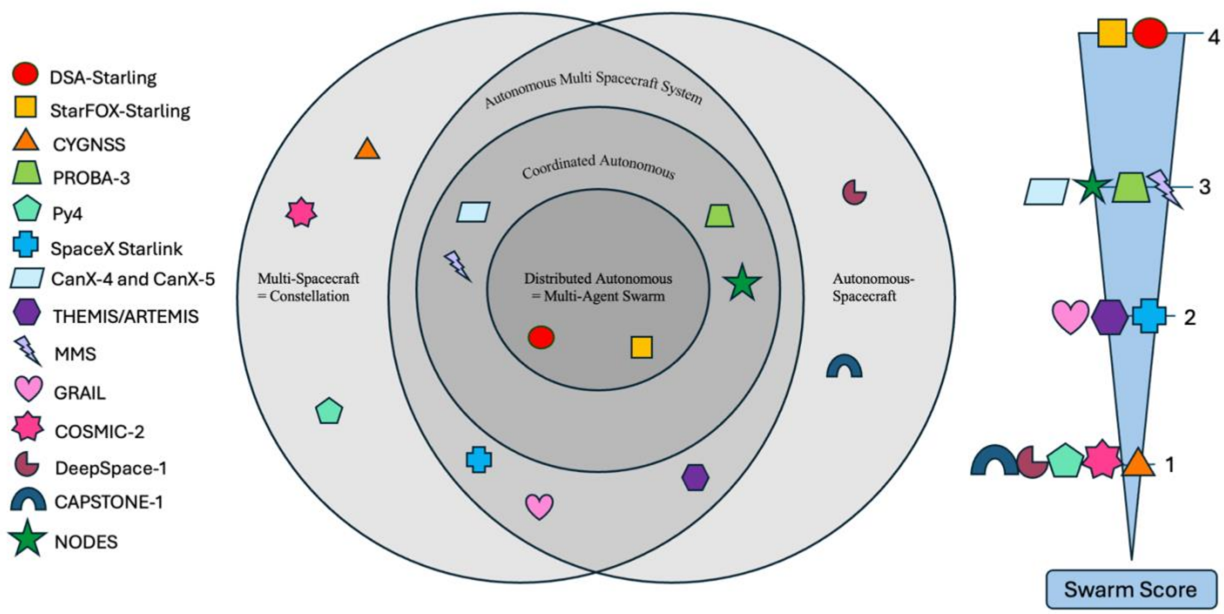

Figure 9: Relevant multi-agent missions that have been successfully demonstrated in orbit. Missions with a higher swarm score demonstrated a greater number of key swarm characteristics, namely autonomy, coordination, and distribution, and thus appear closer to the center of the Venn diagram.

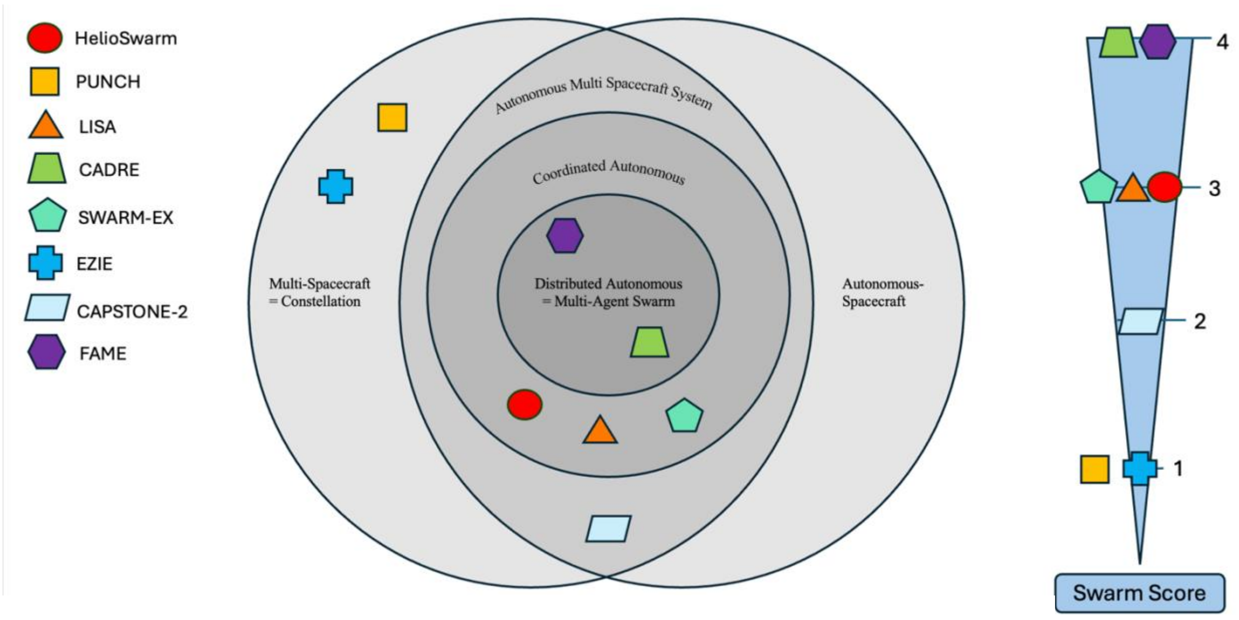

Figure 10: Multi-agent missions that are planned or have yet to demonstrate relevant multi-agent capabilities. Missions with a higher swarm score will have demonstrated a greater number of key swarm characteristics, namely autonomy, coordination, and distribution, and thus appear closer to the center of the Venn diagram.

Appendix

STMD Civil Space Shortfalls

The following are the Civil Space Shortfalls that the Space Technology Mission Directory has identified that relate to the Next Generation Multi-agent Swarm Study.

ID	Title	Description
1430	Small Spacecraft Propulsion	Reliably manufacturable onboard propulsion systems than can impart more ΔV are needed for both nanosatellites (CubeSat) and microsatellites (large CubeSat / small ESPA (EELV Secondary Payload Adapter) class.
1431	Access Beyond LEO for Small Spacecraft	This shortfall addresses the need for vehicles, platforms, and technologies that enable access to non-traditional orbits, including destinations beyond LEO such as Cislunar space, Lagrange points, Martian space, Venusian space, and additional deep space locations.
1433	Position, Navigation, and Timing for Small Spacecraft	Continued small spacecraft use in deep space requires highly accurate position knowledge and precision timing that does not depend on GPS or other Earth-centric aids.
1438	Autonomy, Edge Computation, and Interoperable Networking for Small Spacecraft	Increasing numbers of small spacecraft and small spacecraft constellations require an improvement in the ability to perform autonomous operations in order to: reduce demand on limited ground network assets; enable rapid, on-board responses to time sensitive scenarios (i.e. conjunction response or a dynamic observation event); and enable multiple small spacecraft to operate as a single distributed system to perform advanced observations applicable to space and earth science without ground stations in the loop.
1557	Position, Navigation, and Timing (PNT) for In-Orbit and Surface Applications	Safe navigation for extended durations and distances on planetary surfaces requires an in-situ navigation and timing system.
1589	Space Situational Awareness	Increased space situational awareness, identification, tracking, conjunction prediction, and maneuver coordination will become critical to an expanded LEO economy as well as USG and NASA activities in near earth locations
1625	Intelligent Multi-Agent Constellations for Cooperative Operations	Autonomous cooperative multi-spacecraft system (i.e., constellation) for cost-effective interdependent and distributed measurement/operations)

Catalog of Multi-Agent Space Systems

The NGS team has compiled a catalog of 50 multi-agent space systems relevant to the scope of this study. The catalog includes past, ongoing, and planned missions that demonstrate or enable multi-agent swarm capabilities. A select few single-agent missions that directly contribute enabling technologies or lessons learned to the next generation of multi-agent swarms were also included. Although there exist numerous single-agent missions that advance the state of the art, particularly in lunar and cislunar environments, many were excluded from the catalog, as the focus of this study is on multi-agent systems specifically. The single-agent missions removed from the NGS catalog are listed below.

Notable single-agent missions that were removed from the catalog:

- BioSentinel
- Lunar Flashlight
- Lunar Trailblazer
- Lunar IceCube
- Lunar Reconnaissance Orbiter (LRO)
- Near-Earth Asteroid (NEA) Scout
- EQUULEUS (EQUilibriUm Lunar-Earth point 6U Spacecraft)
- CubeSat Mission to Study Solar Particles (CuSP)
- Team Miles
- ArgoMoon
- Lunar Polar Hydrogen Mapper (LunaH-Map)
- Advanced Composite Solar Sail System (ACS3)

Additionally, for sake of brevity, a number of multi-agent system were also removed from the catalog. These missions were generally excluded because they were not discussed in this report and had only minimal information available for catalog entries. In most cases, the systems were either not directly relevant to the scope of the study or lacked sufficient publicly available data to allow for meaningful comparison. The list provided here should therefore be viewed as a non-comprehensive subset of known multi-agent systems, retained for reference but not directly discussed in this report.

Notable multi-agent systems that were removed from the catalog:

- ViaSat-1
- ViaSat-2
- ViaSat-3
- Thuraya 1-4
- Eutelsat 5 West A
- Eutelsat 5 West B
- Starshield
- Hongyun
- Eutelsat 139 A West
- Eutelsat Konnect
- Guowang
- OneWeb
- ViaSat Anik-F2
- ViaSat KA-SAT
- Disaster Monitoring Constellation (DMC)

Mission	*Mission Objective*	*Tech Demo*	*Max Nodes*	*Crosslink*	*Heterogeneity*	*Target Orbit*	*Mass (kg)*	*Lifetime (y)*	*Autonomous?*	*Coordinated?*	*Distributed?*
CAPSTONE	**Validate spacecraft navigation technologies. Verify the feasibility of Near-Rectilinear Halo Orbits. Precursor to the Gateway lunar outpost.**	**1**	**1**	**N/A**	**N/A**	**Lunar NRHO**	**25**	**2.75**	**1**	**0**	**0**
Deep Space 1	**Engineering test flight for 12 new technologies**	**1**	**1**	**N/A**	**N/A**	**Solar orbit**	**486**	**1**	**1**	**0**	**0**
OPS-SAT	**Provided an in-orbit testbed for software experiments**	**1**	**1**	**N/A**	**N/A**	**LEO**	**7**	**4.5**	**1**	**0**	**0**
MarCO	**First CubeSats to operate beyond LEO, Comms relay for InSight lander, Demonstration of comms technology**	**1**	**2**	**UHF radio receiver and X-band radio transmitter**	**Homogeneous**	**Earth escape**	**14**	**0.58**	**0**	**0**	**0**
NODES	**Swarm network capabilities**	**1**	**2**	**AstroDev Lithium 1 UHF transceiver**	**Homogeneous**	**LEO**	**2**	**0.5**	**1**	**1**	**0**
ESA Swarm	**Study Earth's magnetic field**	**0**	**3**	**N/A, each satellite communicates directly with ground stations via S-band frequencies**	**Homogeneous**	**LEO**	**468**	**11**	**0**	**0**	**0**
Astrobee (Survey Manager)	**Mobile camera, Autonomous surveying, Research hosting**	**1**	**3**	**Distributed, DDS and RAPID, ISS Wi-Fi, and Ku-band**	**Homogeneous**	**N/A**	**10**	**1**	**1**	**1**	**0**
CADRE	**Demonstrate the feasibility of using a team of small, autonomous rovers to collaboratively explore and map the Moon's surface and subsurface without direct human control**	**1**	**3 nodes + base station**	**Communicates via mesh network radio, distributed**	**Homogeneous**	**N/A**	**N/A**	**0.04**	**1**	**1**	**1**

PY4	**Demonstration multi-spacecraft communications**	**1**	**4**	**S-band crosslink**	**Homogeneous**	**LEO**	**unknown**	**unknown**	**0**	**0**	**0**
Starling 1.0	**Demonstrate in-flight autonomy**	**1**	**4**	**Distributed crosslink**	**Homogeneous**	**LEO**	**unknown**	**> 3**	**1**	**1**	**1**
SunRISE	**Study solar activity, observing low radio frequency emissions Informs generation of solar particle storms**	**0**	**6**	**No crosslink**	**Homogeneous**	**GEO**	**120, MEV 92.5 kg plus 30% margin**	**1**	**0**	**0**	**0**
COSMIC-2	**Meteorology, ionosphere, climatology, and space weather research**	**0**	**6**	**No crosslink**	**Homogeneous**	**LEO**	**300**	**5**	**0**	**0**	**0**
HelioSwarm	**Study solar wind**	**1**	**8 nodes + hub**	**Centralized through hub**	**Heterogeneous**	**HEO**	**70**	**200**	**1**	**1**	**0**
EDSN	**Demonstrate satellite networking for intersatellite communications**	**1**	**8**	**AstroDev Lithium 1 UHF transceiver**	**Homogeneous**	**LEO**	**1.7 – 2**	**0.16**	**1**	**1**	**0**
Planet Labs SkySat	**High-resolution Earth observation imagery**	**0**	**21**	**No crosslink**	**Homogeneous**	**LEO, SSO**	**~110**	**6**	**0**	**0**	**0**
Planet Labs Dove	**High-resolution Earth observation imagery**	**0**	**150**	**No crosslink**	**Homogeneous**	**LEO, SSO**	**~5.8**	**3**	**0**	**0**	**0**
Planet Labs RapidEye	**High-resolution Earth observation imagery**	**0**	**5**	**No crosslink**	**Homogeneous**	**LEO, SSO**	**~150 – 156**	**> 11**	**0**	**0**	**0**
Pléiades Neo	**High-resolution Earth observation imagery**	**0**	**4**	**No Crosslink**	**Homogeneous**	**LEO**	**920**	**10**	**0**	**0**	**0**
IRIDIUM	**Global mobile telecommunications services**	**0**	**66**	**L and Ka-band**	**Heterogeneous**	**LEO**	**860**	**15 – 20**	**0**	**0**	**0**
Telesat LEO	**Satellite internet**	**0**	**1,600 planned, 120 operational**	**Ka-band**	**Heterogeneous**	**LEO**	**800**	**10**	**0**	**0**	**0**
Project Kuiper	**Satellite internet**	**0**	**3,000 planned, 2 operational**	**Optical inter-satellite link**	**Homogeneous**	**LEO**	**~537 – 571**	**7**	**0**	**0**	**0**

SpaceX Starlink	**Satellite internet**	**0**	**5,501 – 6,496 operational**	**Distributed crosslink**	**Heterogeneous**	**LEO**	**227 – 1,250**	**5**	**unknown**	**unknown**	**unknown**
Starling-DSA	**One of the core payloads on the NASA Starling mission, demonstrated collaborative resource allocation for multi-point science data collection**	**1**	**4**	**Distributed crosslink**	**Homogeneous**	**LEO**	**unknown**	**> 3**	**1**	**1**	**1**
Starling-StarFOX	**One of the core payloads on the NASA Starling mission, demonstrated target detection, recognition, and tracking using optical imagery**	**1**	**4**	**Distributed crosslink**	**Homogeneous**	**LEO**	**unknown**	**> 3**	**1**	**1**	**1**
Starling-ROMEO	**One of the core payloads on the NASA Starling mission, designed to autonomously maintain a satellite swarm in a desired formation**	**1**	**4**	**Distributed crosslink**	**Homogeneous**	**LEO**	**unknown**	**N/A**	**1**	**1**	**1**
EZIE	**Study auroral electrojets in the upper atmosphere of Earth**	**0**	**3**	**No crosslink**	**Homogeneous**	**LEO**	**unknown**	**1.5**	**0**	**0**	**0**
BRITE	**Study the variability of the brightest and most massive stars in the sky through precision photometry**	**0**	**6**	**No crosslink**	**Heterogeneous**	**LEO**	**~7**	**> 2**	**0**	**0**	**0**
LISA	**Determine how and when massive black holes have formed and grown over cosmic time**	**0**	**3**	**No crosslink**	**Homogeneous**	**Heliocentric**	**unknown**	**6.5**	**1**	**1**	**0**
ESCAPADE	**Investigate the structure, composition, variability, and dynamics of the magnetosphere of Mars**	**0**	**2**	**No crosslink**	**Homogeneous**	**Highly elliptical orbit around Mars**	**200 kg dry mass, < 550 kg wet mass**	**0.9**	**unknown**	**unknown**	**unknown**
CYGNSS	**Collect frequent space-based measurements of surface wind speeds in the inner core of tropical cyclones**	**0**	**8**	**No crosslink**	**Homogeneous**	**LEO**	**~25**	**2**	**0**	**0**	**0**

MagCon	**Study the transport of mass and energy across the magnetosphere of Earth**	**0**	**36**	**No crosslink**	**Homogeneous**	**Highly elliptical, near equatorial orbits**	**59.82 kg dry mass, 86 kg wet mass**	**N/A**	**0**	**0**	**0**
TROPICS	**Perform detailed and frequent studies of tropical cyclones**	**0**	**4**	**No crosslink**	**Homogeneous**	**LEO**	**5.34**	**> 2**	**0**	**0**	**0**
A-Train	**Earth observation constellation**	**0**	**4**	**No crosslink**	**Heterogeneous**	**LEO, SSO**	**N/A**	**N/A**	**0**	**0**	**0**
THEMIS	**Study how mass and energy move through the near-Earth space environment**	**0**	**5**	**No crosslink**	**Homogeneous**	**LEO**	**126-128**	**> 10**	**1**	**0**	**0**
THEMIS-ARTEMIS	**Follow-on to the THEMIS mission, 2 of the 5 satellites were redirected to the Moon to study how mass and energy move through the lunar space environment**	**0**	**2**	**No crosslink**	**Homogeneous**	**Cislunar**	**126-128**	**> 10**	**1**	**0**	**0**
PREFIRE	**Study heat emissions in the Arctic**	**0**	**2**	**No crosslink**	**Homogeneous**	**LEO, SSO**	**3**	**> 1**	**0**	**0**	**0**
PUNCH	**Make global, 3D observations of the entire inner heliosphere to learn how the Sun's corona becomes solar wind**	**0**	**4**	**No crosslink**	**Homogeneous**	**LEO, SSO**	**40**	**2**	**0**	**0**	**0**
RCM	**Provide data for climate research and commercial applications including oil exploration, fishing, shipping, etc.**	**0**	**3**	**No crosslink**	**Homogeneous**	**LEO**	**466.67**	**7**	**0**	**0**	**0**
GRACE-FO	**Map the gravitational field of Earth**	**0**	**2**	**Ka/K-band ranging system**	**Homogeneous**	**LEO**	**487**	**15.6**	**1**	**0**	**0**
GRAIL	**Follow-up to GRACE mission to map the gravitational field of the Moon**	**0**	**2**	**Dual-band ranging system, used Ka-/S-band**	**Homogeneous**	**Low altitude, near circular,**	**132.6 kg dry mass, 202.4 kg wet mass**	**1.25**	**1**	**0**	**0**

						lunar polar orbit					
CLICK-B/C	**Demonstrate laser communications between CubeSats**	**1**	**2**	**Laser comm**	**Homogeneous, but use different laser wavelengths for two-way comm**	**LEO**	**unknown**	**unknown**	**0**	**0**	**0**
DICE	**Measure ionosphere plasma density**	**0**	**2**	**No crosslink**	**Homogeneous**	**LEO, SSO**	**1.42**	**> 0.25**	**0**	**0**	**0**
PROBA-3	**High-precision satellite formation flying**	**1**	**2**	**Radio-based inter-satellite link, S-band**	**Heterogeneous**	**Highly elliptical Earth orbit**	**550 (combined mass in stack)**	**2**	**1**	**1**	**0**
MMS	**Study the magnetosphere of the Earth**	**0**	**4**	**Radio-based inter-satellite link, S-band communication transceiver**	**Homogeneous**		**1,360 (combined mass)**	**> 10**	**1**	**1**	**0**
CanX-4/5	**On-orbit formation flying**	**1**	**2**	**Radio-based inter-satellite link, S-band**	**Homogeneous**	**LEO, SSO**	**7**	**< 0.02**	**1**	**1**	**0**
M2	**Close proximity satellite formation flying and propellant-less formation control**	**1**	**2**	**Radio-based inter-satellite link, S-band**	**Near-identical, share the same physical form factor but different experimental payloads**	**LEO**	**10**	**unknown**	**0**	**0**	**0**
SWARM-EX	**Study the spatial and temporal variability of ion-neutral interactions around Earth**	**0**	**3**	**Inter-satellite crosslink via UHF radios**	**Homogeneous**	**LEO**	**unknown**	**~1**	**1**	**1**	**0**
VISORS	**Perform precision formation flying to detect and study energy release regions of the solar corona**	**1**	**2**	**Inter-satellite crosslink, 5.8 GHz frequency band**	**Heterogeneous**	**LEO, SSO**	**12**	**0.5**	**1**	**1**	**0**

OLFAR	**Build a space-based radio telescope, using a swarm of nanosatellites, to observe the universe at ultra-low radio frequencies**	**unknow n**	**50**	**Inter-satellite crosslink**	**Homogeneous**	**Moon orbit, Earth-Moon L2 or Sun-Earth L4/5**	**N/A**	**unknown**	**1**	**1**	**1**
JPSS	**Provide global environment data used in weather prediction models for forecasts and scientific data used for climate monitoring**	**0**	**5**	**No crosslink**	**Similar, not identical**	**LEO, SSO**	**~2,229 per satellite**	**5 – 7**	**0**	**0**	**0**
GOMX-4	**Demonstrate key technologies to handle large satellite formations**	**1**	**2**	**Inter-satellite crosslink, S-band**	**Heterogeneous, similar platform but different payloads**	**LEO, SSO**	**~8**	**4**	**1**	**1**	**0**

References

Abro, G. E. M., Ali, Z. A., & Masood, R. J. (2024). Synergistic UAV Motion: A Comprehensive Review on Advancing Multi-Agent Coordination. *IECE Transactions on Sensing, Communication, and Control*, *1*(2), 72–88. https://doi.org/10.62762/TSCC.2024.211408

Adams, C. A., Kempa, B. C., Vaughan, W. B., Iatauro, M. J., & Frank, J. D. (2023, August 5). *An Overview of Distributed Spacecraft Autonomy at NASA Ames*. 37th Annual Small Satellite Conference, Logan, UT. https://ntrs.nasa.gov/citations/20230005355

Agarwal, R., Oh, B., Fitzpatrick, D., Buynovskiy, A., Lowe, S., Lisy, C., Kriezis, A., Lan, B., Lee, Z., Thomas, A., Wallace, B., Costantino, E., Miner, G., Thayer, J., D'Amico, S., Lemmer, K., Lohmeyer, W., & Palo, S. (n.d.). *Coordinating Development of the SWARM-EX CubeSat Swarm Across Multiple Institutions*.

Airbus to reshape Earth observation market with its Pleiades Neo constellation. (n.d.). Retrieved June 11, 2025, from https://www.spacedaily.com/reports/Airbus_to_reshape_Earth_observation_market_with_its_Pleiades_Neo_constellation_999.html

Amaro-Seoane, P., Audley, H., Babak, S., Baker, J., Barausse, E., Bender, P., Berti, E., Binetruy, P., Born, M., Bortoluzzi, D., Camp, J., Caprini, C., Cardoso, V., Colpi, M., Conklin, J., Cornish, N., Cutler, C., Danzmann, K., Dolesi, R., … Zweifel, P. (2017). *Laser Interferometer Space Antenna* (No. arXiv:1702.00786). arXiv. https://doi.org/10.48550/arXiv.1702.00786

Angelopoulos, V. (2008). The THEMIS Mission. *Space Science Reviews*, *141*(1), 5–34. https://doi.org/10.1007/s11214-008-9336-1

Arunkumar, E., Paletta, A., & Lightsey, G. (n.d.). *Design of the Hosted Software Application for the VISORS Mission*.

B.A.T.M.A.N. advanced—Open-mesh.org latest documentation. (n.d.). Retrieved September 29, 2025, from https://www.open-mesh.org/doc/batman-adv/Wiki.html

Bester, M., Cosgrove, D., Frey, S., Marchese, J., Burgart, A., Lewis, M., Roberts, B., Thorsness, J., McDonald, J., Pease, D., Picard, G., Eckert, M., & Dumlao, R. (2014). ARTEMIS operations—Experiences and lessons learned. *2014 IEEE Aerospace Conference*, 1–20. https://doi.org/10.1109/AERO.2014.6836399

Blackwell, W. J., Braun, S., Bennartz, R., Velden, C., DeMaria, M., Atlas, R., Dunion, J., Marks, F., Rogers, R., Annane, B., & Leslie, R. V. (2018). An overview of the TROPICS NASA Earth Venture Mission. *Quarterly Journal of the Royal Meteorological Society. Royal Meteorological Society (Great Britain)*, *144*(Suppl Suppl 1), 16–26. https://doi.org/10.1002/qj.3290

Bonin, G., Roth, N., Armitage, S., Newman, J., Risi, B., & Zee, R. E. (n.d.). *CanX–4 and CanX–5 Precision Formation Flight: Mission Accomplished!*

Boshuizen, C., Mason, J., Klupar, P., & Spanhake, S. (2014, August 4). Results from the Planet Labs Flock Constellation. *Small Satellite Conference*. https://digitalcommons.usu.edu/smallsat/2014/PrivEnd/1

Brown, M., Peters, E. G. W., Dilkes, M., Jeffreson, R., Lambert, A., Saleem, R., Boyce, R., Kruzins, E., & Bateman, T. (2023). *Autonomous Close Proximity Differential Drag Control of Satellite Formations using an Inter-Satellite Radio Frequency Link*.

Carreno-Luengo, H., Crespo, J. A., Akbar, R., Bringer, A., Warnock, A., Morris, M., & Ruf, C. (2021). The CYGNSS Mission: On-Going Science Team Investigations. *Remote Sensing*, *13*(9), Article 9. https://doi.org/10.3390/rs13091814

Chai, D. J., Queen, S. Z., & Placanica, S. J. (2015, October 19). *Precision Closed-Loop Orbital Maneuvering System Design and Performance for the Magnetospheric Multiscale Formation*. International Symposium on Space Flight Dynamics, Munich. https://ntrs.nasa.gov/citations/20150020485

Chaudhry, A. U., & Yanikomeroglu, H. (2021a). Laser Intersatellite Links in a Starlink Constellation: A Classification and Analysis. *IEEE Vehicular Technology Magazine*, *16*(2), 48–56. https://doi.org/10.1109/MVT.2021.3063706

Chien, S., Zilberstein, I., Candela, A., Barretta, D., Rijlaarsdamn, D., Hendrix, Tl, Dunne, A., Cortes Grau, O., Gol i Mestre, A., Pedra Bove, (2025).

(2021b). Laser Intersatellite Links in a Starlink Constellation: A Classification and Analysis. *IEEE Vehicular Technology Magazine*, *16*(2), 48–56. https://doi.org/10.1109/MVT.2021.3063706

Cote, S., Lapointe, M., De Lisle, D., Arsenault, E., & Wierus, M. (2021). The RADARSAT Constellation: Mission Overview and Status. *EUSAR 2021; 13th European Conference on Synthetic Aperture Radar*, 1–5. https://ieeexplore.ieee.org/document/9472534/

De La Croix, J.-P., Rossi, F., Brockers, R., Aguilar, D., Albee, K., Boroson, E., Cauligi, A., Delaune, J., Hewitt, R., Kogan, D., Lim, G., Morrell, B., Nakka, Y., Nguyen, V., Proença, P., Rabideau, G., Russino, J., Da Silva, M. S., Zohar, G., & Comandur, S. (2024). Multi-Agent Autonomy for Space Exploration on the CADRE Lunar Technology

Demonstration. *2024 IEEE Aerospace Conference*, 1–14. https://doi.org/10.1109/AERO58975.2024.10521425

Donnellan, A. (2021). *Observing Earths Surface Topography and Vegetation Structure in the Next Decade: Science Objectives and Needs*. *2021*, U32A-01. https://science.nasa.gov/earth-science/decadal-surveys/decadal-stv/

Dorri, A., Kanhere, S. S., & Jurdak, R. (2018). Multi-agent systems: A survey. *IEEE Access*, *6*, 28573–28593.

Doyon, M., Smyth, J., Kroupnik, G., Carrie, C., Sauvageau, M., Levesque, J.-F., Babiker, F., Abbasi, V., Giguere, C., Côté, S., & Bergeron, J. (2018, May 28). RADARSAT CONSTELLATION MISSION: Toward launch and operations. *2018 SpaceOps Conference*. 2018 SpaceOps Conference, Marseille, France. https://doi.org/10.2514/6.2018-2605

Edwards-Stewart, C. (n.d.). *NASA's GRAIL Spacecraft Formation Flight, End of Mission Results, and Small-Satellite Applications*.

Friis-Christensen, E., Lühr, H., Knudsen, D., & Haagmans, R. (2008). *Swarm* – An Earth Observation Mission investigating Geospace. *Advances in Space Research*, *41*(1), 210–216. https://doi.org/10.1016/j.asr.2006.10.008

Gardner, T., Cheetham, B., Bolliger, M., Caudill, M., DeMoudt, L., Forsman, A., Hanf, T., Kam, A., Kayser, E., Ott, C., Re, N., Rogers, R., Zara, A., Umansky, H., Clarke, K., Orbital, T., Bryant, B., & Svitek, T. (n.d.). *CAPSTONE: An Ongoing Demonstration of Navigation and Autonomy Technologies in the Cislunar Domain*.

Gardner, T., Cheetham, B., Parker, J., Forsman, A., Kayser, E., Thompson, M., Ott, C., DeMoudt, L., Caudill, M., Bolliger, M., Kam, A., Thompson, K., Rogers, R., Umansky, H., Bryant, B., & Svitek, T. (n.d.). *CAPSTONE: A Summary of a Highly Successful Mission in the Cislunar Environment*.

Gjerloev, J. W., Yee, J.-H., Kepko, L., Laundal, K., Merkin, V. G., Alken, P., Maute, A. I., Vanhamaki, H., Araujo de Mesquita, R. L., Madelaire, M., Misra, S., Padmanabhan, S., Verkhoglyadova, O. P., Swartz, W. H., Wang, W., Espy, P. J., & Larsson, R. (2024). Electojet Zeeman Imaging Explorer (EZIE): First light. *AGU Fall Meeting Abstracts*, *2024*(2499), SA33A-2499.

GRACE-FO: The Gravity Recovery and Climate Experiment Follow-On Mission. (n.d.). https://doi.org/10.2514/1.A34326

Handler, G., Pigulski, A., Weiss, W. W., Moffat, A. F. J., Kuschnig, R., Wade, G. A., Orleanski, P., Rucinski, S. M., Koudelka, O., Smolec, R., Zwintz, K., Matthews, J. M., Popowicz, A., Baade, D., Neiner, C., Pamyatnykh, A. A., Rowe, J., & Schwarzenberg-Czerny, A. (2017). The BRITE-Constellation Nanosatellite Space Mission And Its First Scientific Results. *EPJ Web of Conferences*, *160*, 01001. https://doi.org/10.1051/epjconf/201716001001

Hanson, J., Allmen, J., Petro, A., & Fishman, J. L. (2014). Nodes-Network & Operation Demonstration Satellite. *Annual AIAA/USU Conference on Small Satellites*.

Holliday, M., Ticknor, R., Stupl, J., Hunter, R., Fisch, P., Sow, I. S., Willis, J., & Manchester, Z. (2024). *The PY4 Mission: A Low-Cost Demonstration of CubeSat Formation-Flying Technologies*.

IEEE Standard for Local and metropolitan area networks—Quality of Service Provision by Network Systems. (2024). *IEEE Std 802.1DC-2024*, 1–54. https://doi.org/10.1109/IEEESTD.2024.10745929

Iorfida, E., Daras, I., Haagmans, R., & Strømme, A. (2023). Swarm A and C Accelerometers: Data Validation and Scientific Interpretation. *Earth and Space Science*, *10*(2), e2022EA002458. https://doi.org/10.1029/2022EA002458

Killough, R., DeForest, C., Seaton, D. B., Smith, K., Hughes, J. M., Lowder, C., Patel, R., & West, M. J. (n.d.). *A Multi-Faceted Approach to Enabling Large-Scale Science in a Microsat Constellation*.

Koenig, A. W., D'Amico, S., & Lightsey, E. G. (2023). Formation Flying Orbit and Control Concept for Virtual Super Optics Reconfigurable Swarm Mission. *Journal of Guidance Control Dynamics*, *46*(9), 1657–1670. https://doi.org/10.2514/1.G007334

Kruger, J., Hwang, S. S., & D'Amico, S. (2024). *Starling Formation-Flying Optical Experiment: Initial Operations and Flight Results* (No. arXiv:2406.06748). arXiv. https://doi.org/10.48550/arXiv.2406.06748

Kruger, J., Koenig, A. W., & D'Amico, S. (2023). Starling Formation-Flying Optical Experiment (StarFOX): System Design and Preflight Verification. *Journal of Spacecraft and Rockets*, *60*(6), 1755–1777. https://doi.org/10.2514/1.A35598

Labrèche, G., Evans, D., Marszk, D., Mladenov, T., Shiradhonkar, V., Soto, T., & Zelenevskiy, V. (2022). OPS-SAT Spacecraft Autonomy with TensorFlow Lite, Unsupervised Learning, and Online Machine Learning. *2022 IEEE Aerospace Conference (AERO)*, 1–17. https://doi.org/10.1109/AERO53065.2022.9843402

Landerer, F. W., Flechtner, F. M., Save, H., Webb, F. H., Bandikova, T., Bertiger, W. I., Bettadpur, S. V., Byun, S. H., Dahle, C., Dobslaw, H., Fahnestock, E., Harvey, N., Kang, Z., Kruizinga, G. L. H., Loomis, B. D., McCullough, C., Murböck, M., Nagel, P., Paik, M., … Yuan, D.-N. (2020). Extending the Global Mass Change Data Record: GRACE

Follow-On Instrument and Science Data Performance. *Geophysical Research Letters*, *47*(12), e2020GL088306. https://doi.org/10.1029/2020GL088306

Le Moigne, J. (2023, February 16). *Novel Observing Strategies for NASA Future Earth Science Missions*. https://ntrs.nasa.gov/citations/20230002165

L'Ecuyer, T. S., Drouin, B. J., Anheuser, J., Grames, M., Henderson, D. S., Huang, X., Kahn, B. H., Kay, J. E., Lim, B. H., Mateling, M., Merrelli, A., Miller, N. B., Padmanabhan, S., Peterson, C., Schlegel, N.-J., White, M. L., & Xie, Y. (2021). The Polar Radiant Energy in the Far Infrared Experiment: A New Perspective on Polar Longwave Energy Exchanges. *Bulletin of the American Meteorological Society*, *102*(7), E1431–E1449. https://doi.org/10.1175/BAMS-D-20-0155.1

Lee, J. D., & See, K. A. (2004). Trust in Automation: Designing for Appropriate Reliance. *Human Factors*, *46*(1), 50–80. https://doi.org/10.1518/hfes.46.1.50_30392

LISA – Laser Interferometer Space Antenna – NASA Home Page. (n.d.). Retrieved August 22, 2025, from https://lisa.nasa.gov/

LISA factsheet. (n.d.). Retrieved September 9, 2025, from https://www.esa.int/Science_Exploration/Space_Science/LISA/LISA_factsheet

Lowe, S., Fitzpatrick, D., Buynovskiy, A., Shoemaker, L., Palo, S., & D'Amico, S. (2024). Concept of Operations for SWARM-EX: A Three CubeSat Formation-Flying Mission. *2024 IEEE Aerospace Conference*, 1–13. https://doi.org/10.1109/AERO58975.2024.10521149

Lowry, A., Bowie, J. T., Munk, M. M., Krenn, A. G., & Guidi, C. (2024). *Stakeholder Webinar: Civil Space Shortfall Ranking: 2024 Feedback Results*. https://ntrs.nasa.gov/citations/20240013282

Lowry, A., & Charania, A. C. (2024). *2024 NASA Technology Taxonomy*. NASA. https://www.nasa.gov/otps/2024-nasa-technology-taxonomy/

M2 Mission | UNSW Canberra Space—UNSW Canberra. (n.d.). UNSW Sites. Retrieved August 26, 2025, from https://www.unsw.edu.au/canberra/our-research/research-centres-institutes/unsw-canberra-space/missions/m2

Mayer, D. J., Fishman, J. L., & Cahoy, K. (2019). CubeSat laser infrared crosslink. *CubeSat Developers Workshop*.

Merancy, N., Grantier, J. A., Bleacher, J. E., Bhakta, S., Goodliff, K. E., Rucker, M. A., Chai, P. R., Chappell, M. W., Morris-Eckart, A., & Cottonham, A. D. (2025, April 1). *Exploration Systems Development Mission Directorate: Moon to Mars Architecture Definition Document, Revision B*. https://ntrs.nasa.gov/citations/20240015571

Miller, S., Adams, C., Alem, N., Cannon, H., Grashuis, R., Hendriks, T., Hwang, S., Iatauro, M., Pires, C., Kruger, J., & others. (2024). *Starling CubeSat Swarm Technology Demonstration Flight Results*.

Miller, S., Adams, C., Alem, N., Cannon, H., Grashuis, R., Hendriks, T., Hwang, S., Pires, C., Field, M., Kruger, J., D'Amico, S., Pi, S., & Mall, L. (n.d.). *Starling CubeSat Swarm Technology Demonstration Flight Results*.

Misra, S., Padmanabhan, S., Kangaslahti, P., Montes, O., Bosch-Lluis, J., Cofield, R., Ramos, I., & Yee, J.-H. (n.d.). *The Electrojet Zeeman Imaging Explorer (EZIE) mission and the Microwave Electrojet Magnetogram (MEM) radiometer instrument*.

NASA's EZIE Launches on Mission to Study Earth's Electrojets—NASA Science. (2025, March 15). https://science.nasa.gov/science-research/heliophysics/nasas-ezie-launches-on-mission-to-study-earths-electrojets/

NASA's PUNCH Mission Reaches Science Orbit, Releases Data—NASA Science. (2025, August 13). https://science.nasa.gov/blogs/punch/2025/08/13/nasas-punch-mission-reaches-science-orbit-releases-data/

Nayak, S., Lim, G., Rossi, F., & Otte, M. (n.d.). *Multi-Robot Exploration for the CADRE Mission*.

OFFensive Swarm-Enabled Tactics. (n.d.). Retrieved June 18, 2025, from https://www.darpa.mil/research/programs/offensive-swarm-enabled-tactics

Pablo, H., Whittaker, G. N., Popowicz, A., Mochnacki, S. M., Kuschnig, R., Grant, C. C., Moffat, A. F. J., Rucinski, S. M., Matthews, J. M., Schwarzenberg-Czerny, A., Handler, G., Weiss, W. W., Baade, D., Wade, G. A., Zocłońska, E., Ramiaramanantsoa, T., Unterberger, M., Zwintz, K., Pigulski, A., … Zee, R. E. (2016). The BRITE Constellation Nanosatellite Mission: Testing, Commissioning, and Operations*. *Publications of the Astronomical Society of the Pacific*, *128*(970), 125001. https://doi.org/10.1088/1538-3873/128/970/125001

Peñin, L. F., Scoarnec, Y., Fernández-Ibarz, J. M., Cazorla, C., Villa, J. R., Kassel, R., Benito, F. J., de Aragón, A., Baudoux, D., Strippoli, L., Galano, D., Mellab, K., Keplerlaan, E., & Galano, D. (n.d.). *Proba-3: ESA's small satellites precise formation flying mission to study the Sun's inner corona as never before*.

Popowicz, A., Pigulski, A., Bernacki, K., Kuschnig, R., Pablo, H., Ramiaramanantsoa, T., Zocłońska, E., Baade, D., Handler, G., Moffat, A. F. J., Wade, G. A., Neiner, C., Rucinski, S. M., Weiss, W. W., Koudelka, O., Orleański, P., Schwarzenberg-Czerny, A., & Zwintz, K. (2017). BRITE Constellation: Data processing and photometry. *Astronomy & Astrophysics*, *605*, A26. https://doi.org/10.1051/0004-6361/201730806

PREFIRE - NASA Science. (2023, June 16). https://science.nasa.gov/mission/prefire/

PROBA-3 (Project for On-Board Autonomy-3)—eoPortal. (n.d.). Retrieved June 18, 2025, from https://www.eoportal.org/satellite-missions/proba-3#proba-3-project-for-on-board-autonomy-3

Project: Space Weather Atmospheric Reconfigurable Multiscale Experiment (SWARM-EX) | Space Rendezvous Laboratory. (n.d.). Retrieved August 26, 2025, from https://slab.stanford.edu/projects/swarmex

PUNCH - NASA Science. (2023, June 15). https://science.nasa.gov/mission/punch/

PUNCH: The Mission. (n.d.). Retrieved August 26, 2025, from https://punch.space.swri.edu/punch_about_mission.php

RADARSAT Constellation—eoPortal. (n.d.). Retrieved June 18, 2025, from https://www.eoportal.org/satellite-missions/rcm#space-and-hardware-components

Satellite Operations and Formation Flying | UNSW Canberra Space—UNSW Canberra. (n.d.). UNSW Sites. Retrieved August 26, 2025, from https://www.unsw.edu.au/canberra/our-research/research-centres-institutes/unsw-canberra-space/our-research/satellite-operations-and-formation-flying

Spence, H. E., Alexandrova, O., Arzamasskiy, L., Argall, M. R., Caprioli, D., Case, A. W., Chandran, B. D. G., Chen, L.-J., Dors, I., Eastwood, J. P., Forsyth, C., Galvin, A. B., Genot, V. N., Halekas, J. S., Hesse, M., Horbury, T. S., Jian, L., Kasper, J. C., Klein, K. G., … Zweibel, E. (2022). *An Overview of HelioSwarm: A NASA MIDEX Mission to Reveal the Nature of Turbulence in Space Plasmas*. *2022*, SH12E-1485. https://ui.adsabs.harvard.edu/abs/2022AGUFMSH12E1485S

Stone, P., & Veloso, M. (2000). Multiagent Systems: A Survey from a Machine Learning Perspective. *Autonomous Robotics*, *8*(3).

Stupl, J. M., & Holliday, M. (2023a). *PY4 Four-CubeSat Swarm of PyCubed-Based Spacecraft*.

Stupl, J. M., & Holliday, M. (2023b). *PY4 Four-CubeSat Swarm of PyCubed-Based Spacecraft*.

Swarm (Geomagnetic LEO Constellation)—eoPortal. (n.d.). Retrieved June 18, 2025, from https://www.eoportal.org/satellite-missions/swarm

Tasking API | Planet Documentation. (2025, May 20). https://docs.planet.com/develop/apis/tasking/

THEMIS (Time History of Events and Macroscale Interactions during Substorms)—eoPortal. (n.d.). Retrieved August 22, 2025, from https://www.eoportal.org/satellite-missions/themis

THEMIS-ARTEMIS - NASA Science. (2017, November 30). https://science.nasa.gov/mission/themis-artemis/

TIRAPLEGUI, S., SERRANO, D., PENIN, L. F., CONTRERAS, R., & RODRIGUEZ, G. (2019). *Proba-3: Challenges and Needs for Sub-Millimetre Autonomous Formation Flying* [PDF]. 15 pages. https://doi.org/10.13009/EUCASS2019-764

Weiss, J.-P., Schreiner, W. S., Braun, J. J., Xia-Serafino, W., & Huang, C.-Y. (2022). COSMIC-2 Mission Summary at Three Years in Orbit. *Atmosphere*, *13*(9), Article 9. https://doi.org/10.3390/atmos13091409

Williams, T. W., Ottenstein, N., Farahmand, M., & Palmer, E. (2016, September 13). Initial Satellite Formation Flight Results from the Magnetospheric Multiscale Mission. *AIAA/AAS Astrodynamics Specialist Conference*. AIAA/AAS Astrodynamics Specialist Conference, Long Beach, California. https://doi.org/10.2514/6.2016-5505

Yost, B., & Weston, S. (2024). *State-of-the-Art Small Spacecraft Technology* (Technical Publication No. NASA/TP-20240001462). NASA Ames Research Center. https://ntrs.nasa.gov/api/citations/20240001462/downloads/2023%20SOA_final.pdf

Zuber, M. T., Smith, D. E., Lehman, D. H., Hoffman, T. L., Asmar, S. W., & Watkins, M. M. (2013). Gravity Recovery and Interior Laboratory (GRAIL): Mapping the Lunar Interior from Crust to Core. *Springer Netherlands*. https://dspace.mit.edu/handle/1721.1/103610