Invited Talk

Title: Priority-Sensitive Mission Continuity for Degraded Multi-Orbit Satellite Constellations

Speaker: Daniel Reynolds, Massachusetts Institute of Technology

Abstract:

Satellite mega-constellations are being deployed to provide resilient global service, yet collisions, severe space weather, and degraded ground coordination threaten the legacy command-and-control assumptions on which those systems depend. This paper evaluates satellite task re-prioritization within a hybrid Mean Field Game Theory (MFGT) and Model Predictive Control (MPC) distributed fallback autonomy framework for heterogeneous multi-orbit constellations under severe degradation. The framework converts task backlog, priority concentration, agent losses, and regional service memory into a broadcast incentive heatmap that re-prioritizes task selection across the surviving constellation, while onboard receding-horizon task selection enforces local feasibility to generate optimal execution threads. In a representative stress event, a 500-satellite, 6-ground-station architecture is reduced to 146 satellites and 2 ground stations before a geographically concentrated 101-task high-priority workload is injected. The result is priority-sensitive capacity redirection under degraded control. All 101 urgent tasks are admitted immediately and completed on time. At injection, 101 of 104 commitments target the urgent workload. Across the run, the framework completes 1,508 of 1,618 tasks, yielding a task completion rate of 93.2% and a priority satisfaction index of 96.3%. These results show degraded-belief fallback autonomy can redirect post-loss capacity fast enough to preserve mission continuity under severe degradation.

Speaker Biography:

Major Daniel Reynolds is an active duty officer in the United States Space Force and a Secretary of the Air Force STEM PhD Fellow in the Department of Aeronautics and Astronautics at the Massachusetts Institute of Technology (MIT), where he is a member of the Space Telecommunications, Astronomy, and Radiation (STAR) Laboratory under the guidance of Prof. Kerri Cahoy. His research focuses on distributed autonomy for resilient space operations, with an emphasis on scalable coordination and control frameworks for heterogeneous satellite constellations operating across multiple orbits under degraded and disrupted conditions. By combining game theory, predictive control, and autonomous decision-making methods, he is working to develop architectures that enable satellites to continue making locally executable, mission-aware decisions even when centralized coordination is limited or unavailable.

Daniel has nearly a decade of experience spanning academic, nonprofit, and U.S. government-service space initiatives. He earned his bachelor's degree in Astronautical Engineering from the United States Air Force Academy in 2017 and his master's degree in Aeronautics and Astronautics from the Massachusetts Institute of Technology in 2019. In 2021, he completed the U.S. Air Force Test Pilot School's Space Test Course, which provided hands-on training in flight-test fundamentals, advanced space system testing, and space operations. Taken together, his academic, operational, and cross-sector experiences inform a research approach centered on operational realism, rigorous system-level evaluation, and the design of autonomy architectures for demanding mission environments.

Priority-Sensitive Mission Continuity for Degraded Multi-Orbit Satellite Constellations

Daniel Reynolds
Massachusetts Institute of Technology
Cambridge, United States
dcr@mit.edu

Kerri Cahoy
Massachusetts Institute of Technology
Cambridge, United States
kcahoy@mit.edu

Olivier de Weck
Massachusetts Institute of Technology
Cambridge, United States
deweck@mit.edu

ABSTRACT

Satellite mega-constellations are being deployed to provide resilient global service, yet collisions, severe space weather, and degraded ground coordination threaten the legacy command-and-control assumptions on which those systems depend. This paper evaluates satellite task re-prioritization within a hybrid Mean Field Game Theory (MFGT) and Model Predictive Control (MPC) distributed fallback autonomy framework for heterogeneous multi-orbit constellations under severe degradation. The framework converts task backlog, priority concentration, agent losses, and regional service memory into a broadcast incentive heatmap that re-prioritizes task selection across the surviving constellation, while onboard receding-horizon task selection enforces local feasibility to generate optimal execution threads. In a representative stress event, a 500-satellite, 6-ground-station architecture is reduced to 146 satellites and 2 ground stations before a geographically concentrated 101-task high-priority workload is injected. The result is priority-sensitive capacity redirection under degraded control. All 101 urgent tasks are admitted immediately and completed on time. At injection, 101 of 104 commitments target the urgent workload. Across the run, the framework completes 1,508 of 1,618 tasks, yielding a task completion rate of 93.2% and a priority satisfaction index of 96.3%. These results show degraded-belief fallback autonomy can redirect post-loss capacity fast enough to preserve mission continuity under severe degradation.

KEYWORDS

Multi-Agent Systems, Satellite Mega-Constellations, Distributed Fallback Autonomy, Mean Field Game Theory, Model Predictive Control, Degraded Operations, Priority-Sensitive Task Re-Prioritization, Resilient Space Systems

ACM Reference Format:
Daniel Reynolds, Kerri Cahoy, and Olivier de Weck. 2026. Priority-Sensitive Mission Continuity for Degraded Multi-Orbit Satellite Constellations. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 10 pages.

1 INTRODUCTION

Earth orbit is becoming a more congested, fragile, and operationally consequential domain. As of October 2025, more than 54,000 trackable objects larger than 10 cm orbit Earth alongside an estimated 1.2 million debris fragments larger than 1 cm [1]. The European Space Agency warns that, when current launch traffic, fragmentation behavior, and disposal performance are projected forward, long-term collision risk in low Earth orbit reaches roughly four times the agency's orbital sustainability threshold, a benchmark representing an acceptable risk level for a stable long-term environment [1]. The Outer Space Institute's CRASH Clock gives that risk operational immediacy, estimating on 20 March 2026 that, in a severe loss-of-control scenario, a catastrophic collision in LEO could occur within just three days [43]. Kinetic hazards are not the only system-level threat: a Carrington-class geomagnetic storm remains a statistically credible event that could degrade or disable satellites and communications infrastructure across multiple orbital regimes [2, 3]. Against this increasingly hazardous backdrop, mega-constellations of hundreds to thousands of spacecraft are rapidly becoming the dominant architectural paradigm for space-based communications services [4–6], promising persistent global coverage amidst a greater global dependence on an increasingly unpredictable space domain. The implication for mega-constellation command and control is stark: future space systems must be prepared to operate under conditions in which asset loss, stale global state, fragmented communications, and urgent regional demand are not edge cases, but expected stressors.

Constellation command and control (C2) can be viewed along a progression from centralized, to decentralized, to distributed architectures, as illustrated in Figure 1. In centralized C2, planning, scheduling, monitoring, and reassignment are concentrated in a primary ground authority, with spacecraft acting primarily as execution nodes for ground-generated commands [9, 10, 13, 46, 58]. This model provides coherent global oversight, but becomes increasingly brittle as constellation size, task tempo, contact intermittency, and disruption severity grow, because mission execution remains dependent on timely ground coordination and refreshed global state [9, 45, 58]. Decentralized C2 moves beyond this single-node model by distributing operational authority across multiple ground stations, gateways, or regional control nodes, reducing bottlenecks and better matching the parallelism of proliferated mega-constellation operations [11, 13–15, 23, 26, 36, 42, 51]. However, even decentralized ground-centric control can remain vulnerable when communication paths fragment, global state becomes stale, or terrestrial coordination layers are themselves degraded. Distributed C2 therefore represents the critical fallback autonomy step: it embeds tactical decision-making and adaptation within the constellation, allowing spacecraft to operate as autonomous agents under partial observability, intermittent communication, dynamic membership, and degraded coordination [8, 13, 23, 24, 47, 48, 56]. The unresolved issue for fallback operations is therefore not simply whether decision authority can move away from a central ground node, but whether a degraded constellation can still align local

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus.

onboard task choices with mission urgency after the coordination substrate itself has degraded.

This paper addresses that fallback-C2 problem by developing and evaluating a distributed autonomy framework for priority-sensitive task selection in a degraded heterogeneous multi-orbit constellation. The fundamental objective is to preserve mission continuity by enabling surviving satellites to recognize where urgent demand is concentrated, redirect limited capacity accordingly, and continue making feasible onboard tasking decisions when ground-mediated coordination is no longer sufficient. The paper makes two focused contributions.

Contributions: First, this work introduces a heatmap-driven distributed task-selection mechanism that converts regional backlog, agent scarcity, task priority, and recent service effectiveness into a broadcast coordination field for guiding local satellite decisions under degraded belief. Second, this work demonstrates adaptive high-priority re-prioritization after severe architectural loss, showing that a survivor constellation can redirect task commitments toward an urgent regional workload after most satellites and ground stations have been removed. Together, these contributions frame distributed fallback autonomy as a direct answer to the future operating environment of space systems, described by the ability to identify what tasks matter most, shift scarce capacity toward it, and continue executing the mission when legacy C2 assumptions fail.

2 BACKGROUND

Distributed fallback autonomy for satellite mega-constellations is not a single-property problem, but a coupled architectural challenge at the intersection of five domains: An operationally useful architecture must scale to large constellations without centralized optimization, dense pairwise coordination, or all-to-all exchange becoming the dominant bottleneck, as motivated by large-population and mean-field formulations [30, 37, 38, 41, 52]. It must support heterogeneous spacecraft, because sensing, relay, communications, storage, power, and service roles are not uniform across operational fleets [12, 21, 35, 39, 40]. It must accommodate multi-orbit structure, since LEO, MEO, and GEO assets differ in geometry, latency, persistence, coverage, and relay utility [19–21, 27, 52]. It must remain useful under degraded operations, including communication loss, partial observability, fragmentation, and stale global context [7, 31, 54–56]. Finally, it must support bounded isolated operations, allowing agents to continue locally rational behavior when timely access to global coordination signals is interrupted [17, 25, 63].

2.1 Scalability

Mean-field and aggregate-population methods provide one path to tractability by replacing multi-agent interactions with low-dimensional population descriptors or measure flows [29, 30, 37, 38, 41]. Rather than requiring each agent to reason over every other agent, these formulations couple local decisions to aggregate state, reducing the coordination burden associated with large interacting populations [18, 30, 37, 41, 52].

This scalability argument becomes operationally important after capacity loss, when the system must rapidly express where remaining capacity is most needed without solving a centralized reassignment problem over all agents and tasks. Prior work extends mean-field and aggregate coordination ideas toward engineering, partial-information, and networked-control settings [16, 31, 49, 50, 55, 56, 62, 64]. Other approaches reduce complexity through structural sparsification, tier-level control variables, local interference neighborhoods, or parallelized equilibrium search [21, 39, 60]. These results support the premise that large-population coordination can be compressed into aggregate signals, including signals that may be imperfect, delayed, or partially observed. However, aggregate coordination alone does not guarantee that a specific spacecraft can execute a specific mission task under timing, access, queue, power, storage, and communication constraints. That gap motivates the mechanism studied here: a degraded-belief regional coordination signal must be paired with local feasibility-constrained task selection.

2.2 Heterogeneity and Multi-Orbit Regions

Operational constellations are not composed of interchangeable agents. Satellites may differ in sensing roles, relay functions, communications payloads, storage capacity, power state, downlink capability, and resilience under disruption. The reviewed literature introduces heterogeneity in several forms, including heterogeneous task graphs and offloading decisions [12], multi-tier LEO network control [21], satellite/UAV-assisted edge computing [40], hybrid-constellation spectrum sharing [39], and satellite-enabled learning or selection problems [35]. Graph-based and coalition-based resource coordination likewise show that network relationships and cooperative structure matter for distributed decision-making [19, 20, 59].

Multi-orbit structure further constrains which agents can service which tasks. In the present architecture, LEO, MEO, and GEO assets play differentiated tactical, regional-relay, and global-support roles; more generally, prior work treats orbital structure through differences in geometry, visibility, routing opportunity, contact windows, and relay utility [19–21, 27, 52, 57]. These formulations show that orbital layer and network geometry shape coordination opportunity. A degraded survivor constellation may retain some regions, paths, or task classes better than others, so post-loss capacity is not merely reduced; it is redistributed unevenly across altitude, coverage, relay access, and service opportunity.

2.3 Degraded and Isolated Operations

Distributed autonomy is most valuable precisely when the coordination substrate is damaged, delayed, or incomplete. The reviewed literature includes several ingredients relevant to this regime: adversarial or unreliable topology in dynamic connectivity games [7], spoofing or corrupted topology discovery in graph-based satellite coordination [19, 20], intermittent contact and storage-limited forwarding in time-expanded routing [27], peer relay and caching under poor channel conditions [33], imperfect monitoring in repeated games [54], and partial-observation or disturbance-rejection treatments in mean-field settings [31, 55, 56]. Local feasibility under communication limits is also supported in decentralized navigation and MPC-oriented settings [17, 25, 61].

These works provide important degraded-operation ingredients, but the present setting emphasizes a more specific operational mode: the constellation has already lost substantial capacity, the

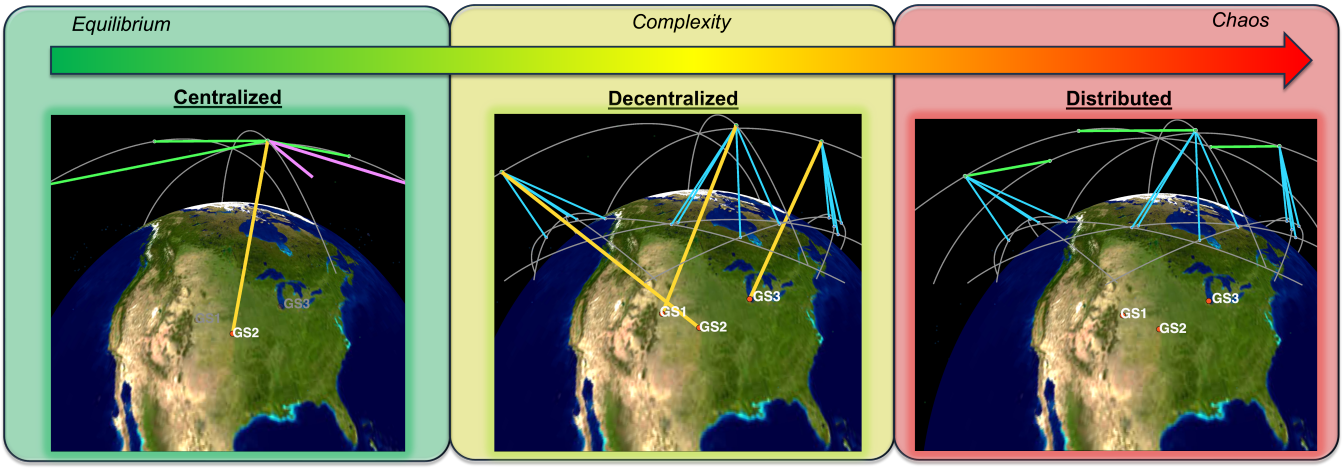

Figure 1: Progression of satellite coordination architectures from centralized to decentralized and distributed control.

ground segment is degraded, and urgent regional work appears after the disruption. Under those conditions, the coordination signal may be stale, delayed, fragmented, or available only through intermittent refresh. Many distributed approaches still depend on iterative exchange, negotiation, reward feedback, neighbor updates, or reliable broadcast information to preserve coordination quality [22, 35, 39, 49, 50, 53, 57, 63]. The unresolved issue is therefore not simply robustness to noisy information. It is whether degraded agents can still redirect scarce capacity toward urgent mission demand when the information needed for coordination is itself imperfect.

Isolation is the limiting case of degraded command and control: an agent may lose timely access to the current constellation state, may be unable to refresh the coordination field, and may still need to choose a feasible action under mission, timing, power, storage, and communication constraints. Prior work provides examples of communication-limited local control, bounded-neighbor reasoning, decentralized response under imperfect monitoring, and transitions away from hub-mediated coordination after link degradation [17, 25, 54, 63]. These works establish that meaningful local behavior can persist when global exchange weakens.

3 FALLBACK AUTONOMY FRAMEWORK

The framework developed in this paper exists at the intersection of the five operational domains identified above through a hybrid Mean Field Game Theory (MFGT) and Model Predictive Control (MPC) architecture. MFGT addresses scalability by compressing large-population interaction into a regional tasking heatmap, while the heatmap-based representation captures demand, scarcity, priority, and recent service history across a heterogeneous multi-orbit operating area. MPC addresses local executability by converting that coordination pressure into onboard task-selection decisions that respect orbit-dependent geometry, timing, queueing, power, storage, and connectivity constraints. Under degraded or partially isolated conditions, satellites act from perceived or cached coordination context rather than relying on centralized reassignment or dense pairwise negotiation. Together, the two layers form a distributed fallback autonomy architecture: a low-bandwidth coordination field identifies where scarce capacity should flow, while each satellite independently determines what it can actually execute.

Mean Field Game Theory (MFGT) provides a scalable way to reason about decision-making in very large multi-agent systems. Rather than modeling every pairwise interaction among agents, mean-field methods approximate the collective influence of the population through an aggregate state, distribution, or measure flow; each agent then computes its response to that aggregate population behavior rather than to every other individual agent [29, 30, 37, 38, 41]. This abstraction preserves the core strategic coupling of a large decision population while avoiding the coordination burden that would arise from dense all-to-all interaction [18, 30, 37, 41, 52]. That structure is well matched to degraded satellite mega-constellation operations: the constellation is large, coupling is driven by regional demand and aggregate agent scarcity, and centralized full-state optimization becomes progressively less tenable as the system scales or degrades [16, 31, 49, 50, 55, 56, 62, 64]. In this work, the mean-field layer implements that idea as a regional broadcast signal: task backlog, agent scarcity, task priority, and recent service effectiveness are compressed into a low-dimensional heatmap that guides local satellite decisions without centralized reassignment, persistent all-to-all coordination, or dense pairwise negotiation.

MFGT alone, however, does not determine whether a specific satellite can physically execute a specific task under degraded conditions. Prior mean-field work generally emphasizes aggregate coordination more directly than task-level execution feasibility, leaving timing, storage, power, queueing, and serviceability constraints to be handled by an additional local control layer [28, 31, 49, 50, 64]. Model Predictive Control (MPC) provides that execution layer. MPC is a receding-horizon control framework in which an agent repeatedly solves a finite-horizon decision problem from its current state, evaluates the predicted consequences of candidate actions, applies only the first selected action, and then replans when the next state is observed [32, 34, 44, 61]. This structure is useful for fallback

autonomy because onboard tasking is not a one-time assignment problem: each satellite must continually decide whether to idle, continue current work, or commit to a new task as orbital access, task deadlines, battery state, storage state, queue load, connectivity, and degraded information evolve [17, 25, 61]. In this work, the MPC layer converts the mean-field heatmap into executable local task-selection decisions by filtering infeasible tasks, scoring admissible actions over a finite horizon, executing the first action, and repeating the process at the next decision epoch. Prior work has shown that mean-field coordination and receding-horizon control can be interfaced in a meaningful way, although not for the full heterogeneous, multi-orbit, degraded-operations constellation setting addressed here [32, 34, 44, 61].

3.1 Mean Field Game Theory

The executable MFGT layer is implemented as a discrete, tiled, simulation-first coordination mechanism rather than as a continuous PDE solver. At each decision epoch, task demand and satellite availability are aggregated over geographic tiles. Each region r is represented by a four-channel broadcast state,

$$\mu_r(t) = \begin{bmatrix} \rho_r(t) \\ \xi_r(t) \\ \delta_r(t) \\ \psi_r(t) \end{bmatrix}, \tag{1}$$

where $\rho_r(t)$ is backlog density, $\xi_r(t)$ is agent scarcity, $\delta_r(t)$ is priority concentration, and $\psi_r(t)$ is service-effectiveness memory. These four terms are designed to capture the mission-relevant conditions that should attract limited surviving capacity during degraded operations.

Backlog density is computed from the number of outstanding, not-completed tasks in region r, normalized by the relative tile-area scale,

$$\rho_r(t) = \frac{|\mathcal{T}_r^{\text{out}}(t)|}{\alpha_r}. \tag{2}$$

Agent scarcity is computed from the number of active satellites that can currently service the region,

$$n_r(t) = \sum_{i \in \mathcal{A}(t)} \mathbf{1}\{i \rightsquigarrow r \text{ at } t\}, \qquad \xi_r(t) = \text{clip}_{[0,1]}\left(\frac{\eta^{\max} - n_r(t)/\alpha_r}{\eta^{\max}}\right). \tag{3}$$

Priority concentration is the mean normalized priority of outstanding tasks in the region,

$$\delta_r(t) = \begin{cases} \dfrac{1}{|\mathcal{T}_r^{\text{out}}(t)|} \displaystyle\sum_{\tau \in \mathcal{T}_r^{\text{out}}(t)} \pi_\tau(t), & |\mathcal{T}_r^{\text{out}}(t)| > 0, \\ 0, & |\mathcal{T}_r^{\text{out}}(t)| = 0, \end{cases} \tag{4}$$

where $\pi_\tau(t) \in [0, 1]$. Finally, service-effectiveness memory is updated as an exponential moving average,

$$\psi_r(t) = (1 - \beta_\psi)\, \psi_r(t - \Delta t) + \beta_\psi\, s_r(t - \Delta t), \tag{5}$$

where $s_r(t - \Delta t)$ is the recent regional completion-fraction proxy. This term prevents the broadcast field from reacting only to instantaneous backlog and allows recent service history to influence regional incentives.

The scalar broadcast incentive for each region is then

$$\Phi_r(t) = \gamma^\top \mu_r(t), \tag{6}$$

where $\gamma \in \mathbb{R}^4$ is a fixed weight vector, meaning that the coordination field is driven entirely by the observed regional mission state. Operationally, $\Phi_r(t)$ functions as a tasking heatmap: regions with high unmet demand, limited available service, high-priority tasks, or poor recent service exert stronger attraction on the surviving constellation.

Because degraded operations may interrupt access to the current broadcast field, each satellite acts on a perceived incentive rather than the true global field. Satellite i maintains a local estimate $\hat{\mu}_{i,r}(t)$, refreshed when communications permit and otherwise retained from cached information. Its perceived regional incentive is

$$\hat{\Phi}_{i,r}(t) = \gamma^\top \mathcal{T}\big(\hat{\mu}_{i,r}(t)\big) + \lambda_r(t), \tag{7}$$

where $\mathcal{T}(\cdot)$ denotes the policy-facing feature transform used by the implementation. This perceived incentive is the only population-level coordination signal used by the local controller. Physical feasibility remains enforced onboard.

3.2 Model Predictive Control

The MPC layer converts the perceived regional incentive field into executable local task-selection decisions. At each epoch, satellite i observes its local state, including orbital access, battery, storage, queue load, resilience state, and connectivity posture. It then constructs a locally feasible candidate task set,

$$\begin{aligned} \mathcal{K}_i^{\text{feas}}(t) = \Big\{ k \in \mathcal{K}(t) \;:\; & \text{svc}_i(k, t) = 1, \\ & t \geq t_k^{\min}, \\ & t + d_k(t) \leq t_k^{\max}, \\ & x_i^{\text{queue}}(t) < Q_i^{\max}, \\ & x_i^{\text{bat}}(t) \geq e_k, \\ & x_i^{\text{stor}}(t) + m_k \leq S_i^{\max} \Big\}. \end{aligned} \tag{8}$$

This set encodes the core execution constraints: serviceability, timing, queue capacity, energy, and storage. Additional implementation guards enforce task-consistency, same-tick claim resolution, and degraded-mode resilience checks.

The admissible action set is composed of idle, continue, and commit actions,

$$\mathcal{U}_i(t) = \{\text{idle}\} \cup \{\text{continue current task}\} \cup \{\text{commit to } k \in \mathcal{K}_i^{\text{feas}}(t)\}. \tag{9}$$

The continue action is available when the satellite has an active queue-head task, while idle remains admissible as a fallback when new commitments are infeasible or unattractive.

For feasible commit actions, the local planner evaluates a mission-structured cost that first scores individual task commitments and then maps those scores into the finite-horizon rollout objective. The resulting stage-cost terms balance resource burden, urgency, latency, regional congestion, capability alignment, and task priority. The mean-field incentive enters this cost through the perceived regional field $\hat{\Phi}_{i,r_k}(t)$, so that tasks located in regions with stronger coordination pressure become more attractive to the onboard planner. Thus, the broadcast heatmap does not assign tasks directly; it biases the local finite-horizon decision problem.

Each satellite then solves a discrete receding-horizon problem,

$$\pi_i^\star(t) \in \arg\min_{\pi \in \Pi_i(t)} \sum_{h=0}^{H-1} \ell_i\Big(\tilde{x}_i(t+h\Delta t),\, \pi(h),\, \hat{\Phi}_{i,\cdot}(t+h\Delta t)\Big), \quad (10)$$

where $\Pi_i(t)$ is the set of admissible horizon-length action sequences, $\tilde{x}_i(\cdot)$ is the rollout state under the planner's internal progression model, and $\ell_i(\cdot)$ is the mission-structured stage cost for idle, continue, or commit actions. In the implemented controller, this stage cost is parameterized by tuned weights on resource preservation, urgency, latency, congestion, alignment, and priority. Only the first action is applied,

$$u_i^\star(t) = \pi_i^\star(t; 0), \quad (11)$$

and the optimization is repeated at the next decision epoch.

This receding-horizon structure allows each satellite to respond to new tasks, changing resource states, evolving orbital access, and degraded broadcast information without requiring centralized reassignment. The implemented controller is therefore not a generic quadratic MPC regulator; it is a mission-structured finite-horizon task planner. The MFGT layer supplies regional coordination pressure through $\hat{\Phi}_{i,r}(t)$, while the MPC layer determines whether acting on that pressure is locally feasible and mission-beneficial.

4 RESULTS AND ANALYSIS

The simulation spans 288 one-minute decision epochs, or approximately 4.8 hours, over a 15×10 southeastern United States (SEUS) regional grid. The initial architecture contains a heterogeneous 500-satellite constellation composed of 400 LEO, 96 MEO, and 4 GEO satellites, supported by six ground stations and initialized with 656 tasks split evenly between remote-sensing and communications demand. During nominal operations, three new tasks are injected per epoch, corresponding to one task-injection opportunity every 60 seconds. At step 96, or t = 1:36:00, a solar-storm-driven mass-casualty event removes 354 satellites, approximately 71% of the constellation, and four ground stations, approximately 67% of the ground segment. The surviving architecture therefore contains 146 satellites and two active ground stations, with the remaining space segment composed of 120 LEO, 24 MEO, and 2 GEO satellites. Eleven minutes later, at step 107, or t = 1:47:00, a geographically concentrated hurricane-response workload injects 101 high-priority tasks across eight coastal tiles in the Florida Atlantic corridor. This event sequence creates a compound stress test: severe architecture loss first reduces available capacity and ground access, then urgent regional demand appears after the constellation is already operating in a degraded state. All controller settings are held fixed across the run, so the observed response reflects closed-loop adaptation rather than phase-specific retuning. The broadcast field is weighted to emphasize priority concentration, backlog pressure, and agent scarcity, while retaining a smaller service-memory term to smooth short-horizon regional response. The onboard receding-horizon planner is tuned to favor urgent and high-priority commitments while still penalizing resource burden, latency, congestion, and poor capability alignment.

4.1 Heatmap-Driven Distributed Task Selection

Figure 3 shows the mechanism behind the first contribution: task selection is driven by a regional incentive field rather than by a centralized reassignment list. The MFGT layer compresses outstanding task density, agent scarcity, priority concentration, and recent service memory into the scalar field $\Phi_r(t)$. Each satellite then sees that field through its local, possibly stale estimate and passes the resulting regional incentive into its onboard receding-horizon task planner. The heatmap effectively serves as the coordination object that changes which feasible commitments look mission-beneficial to locally planning satellites.

The four event-aligned stills show the field responding to the actual mission state. At the initial epoch, the field is broad because the seeded workload is large and widely distributed. By the pre-constellation loss frame, the nominal controller has mostly absorbed that initial backlog, and the field is correspondingly lower and more localized. The step-96 constellation loss changes the interpretation of the remaining work: even with little pending demand, the loss of servicing capacity and ground infrastructure makes regional scarcity more consequential. At step 107, the high-priority injection creates a concentrated ridge over the Florida Atlantic corridor, making the urgent region the dominant attractor for survivor capacity.

The trace metrics show that this visual field shift is paired with a real change in task-selection behavior. The architecture does not preserve nominal volume after the loss event: completed-task flow falls from $860/96 = 8.96$ tasks per step in the nominal phase to $648/192 = 3.38$ tasks per step under constellation loss, and accepted commitments fall from 9.7812 to 3.9427 per step. Queue load contracts as well, with mean queued active pairs dropping from 220.49 before loss to 100.71 afterward. The test therefore exposes the controller to a genuine reduction in available execution capacity before the urgent workload appears.

Within that lower-capacity regime, however, the value signal is preserved more strongly than the raw task count. Priority-weighted completion throughput decreases from 0.9086 in the nominal phase to 0.7944 after loss, then rises to 0.8257 during the high-priority phase. This pattern is the central evidence for heatmap-driven task selection: the post-loss constellation cannot recover nominal task volume, but the broadcast field redirects the remaining feasible commitments toward higher-value demand. At the injection epoch, 101 of 104 accepted commitments are high-priority tasks, showing that the field concentration over the corridor is converted into local task commitments rather than remaining only a visualization artifact.

4.2 Adaptive High-Priority Re-Prioritization After Severe Loss

Figure 4 shows the closed-loop task-flow consequence of the heatmap shift. The created, assigned, and completed traces first show nominal absorption of the seeded workload. At step 96, the constellation is reduced from 500 to 146 satellites and from six to two active ground stations, producing a visible assignment response as the survivor architecture re-forms its queues. Eleven minutes later, the step-107 high-priority injection produces a much sharper selection

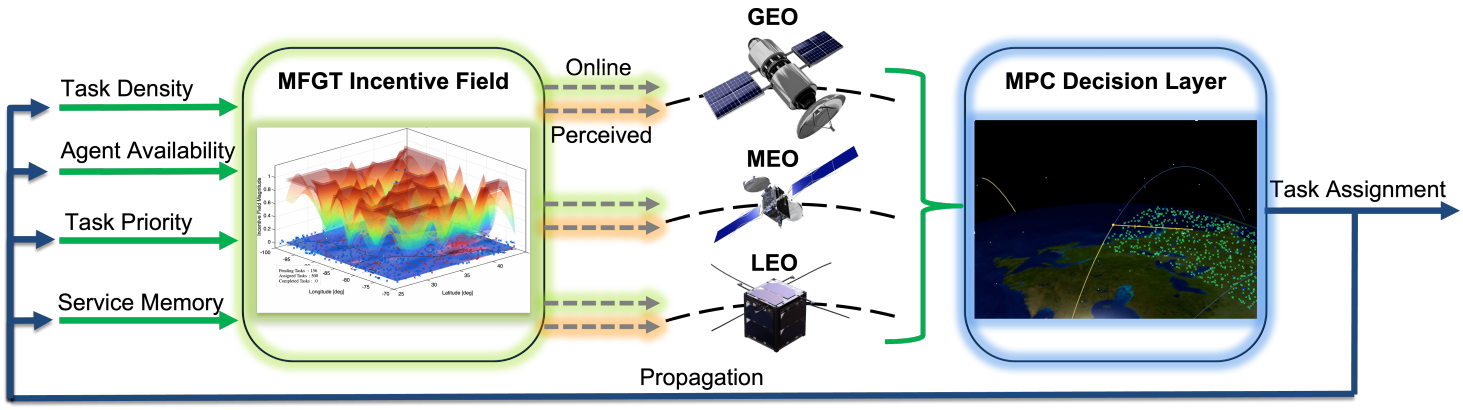

Figure 2: Framework overview showing how regional demand, availability, priority, and service memory form an MFGT incentive field that guides multi-orbit MPC task assignment.

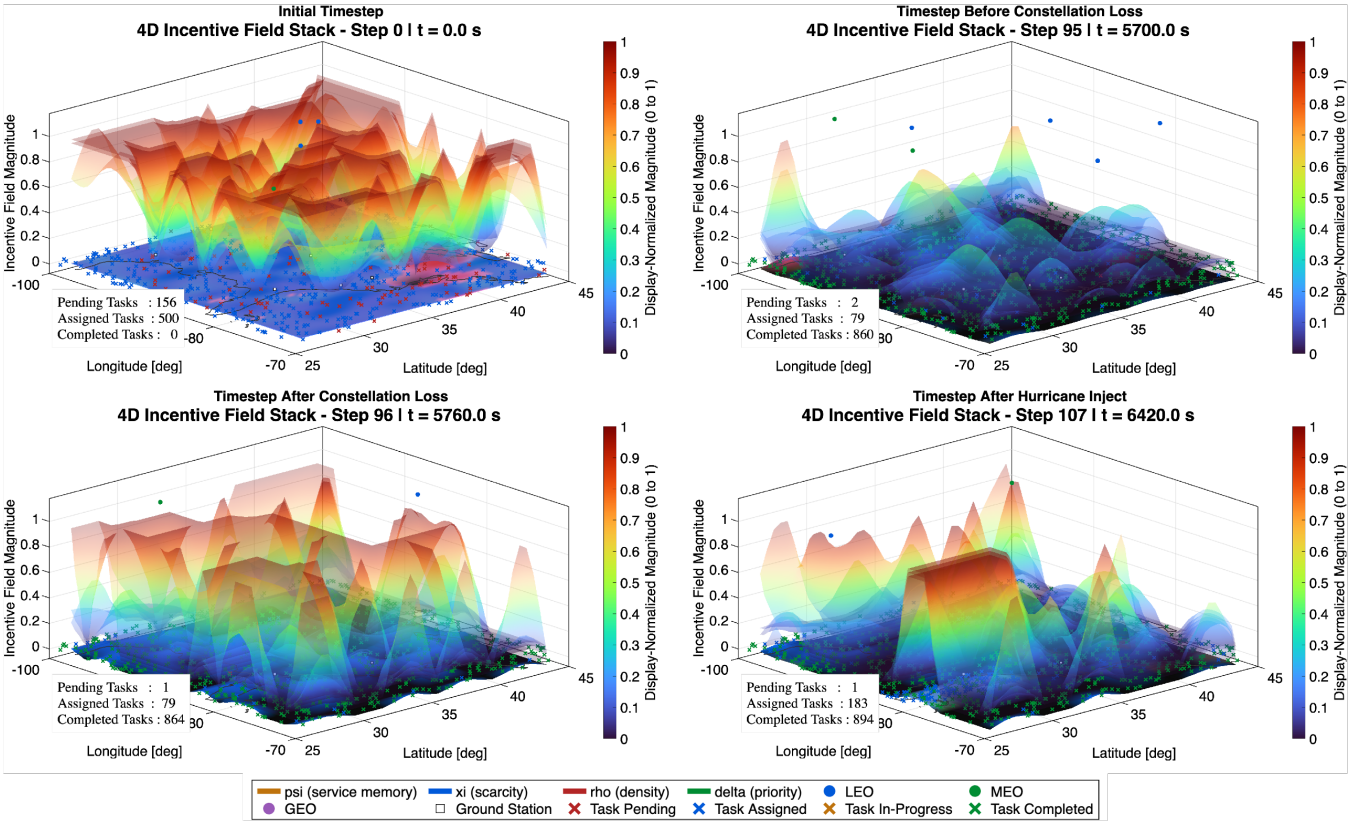

Figure 3: Event-aligned incentive-field stills showing the evolution of regional coordination pressure before contraction, after contraction, and after high-priority task injection.

event: 104 commitments are accepted at that epoch, and 101 of them are directed to the injected urgent tasks.

The high-priority response is immediate at assignment time and bounded at completion time. All 101 injected tasks are assigned at step 107, giving zero-step assignment latency at the mean, median, and 95th percentile. All 101 are also completed on time. Completion latency is tightly concentrated, with mean, median, and 95th-percentile values of 13.60, 14.00, and 15.00 steps, corresponding to approximately 816, 840, and 900 seconds at the one-minute decision cadence. The completion wave visible after the injection

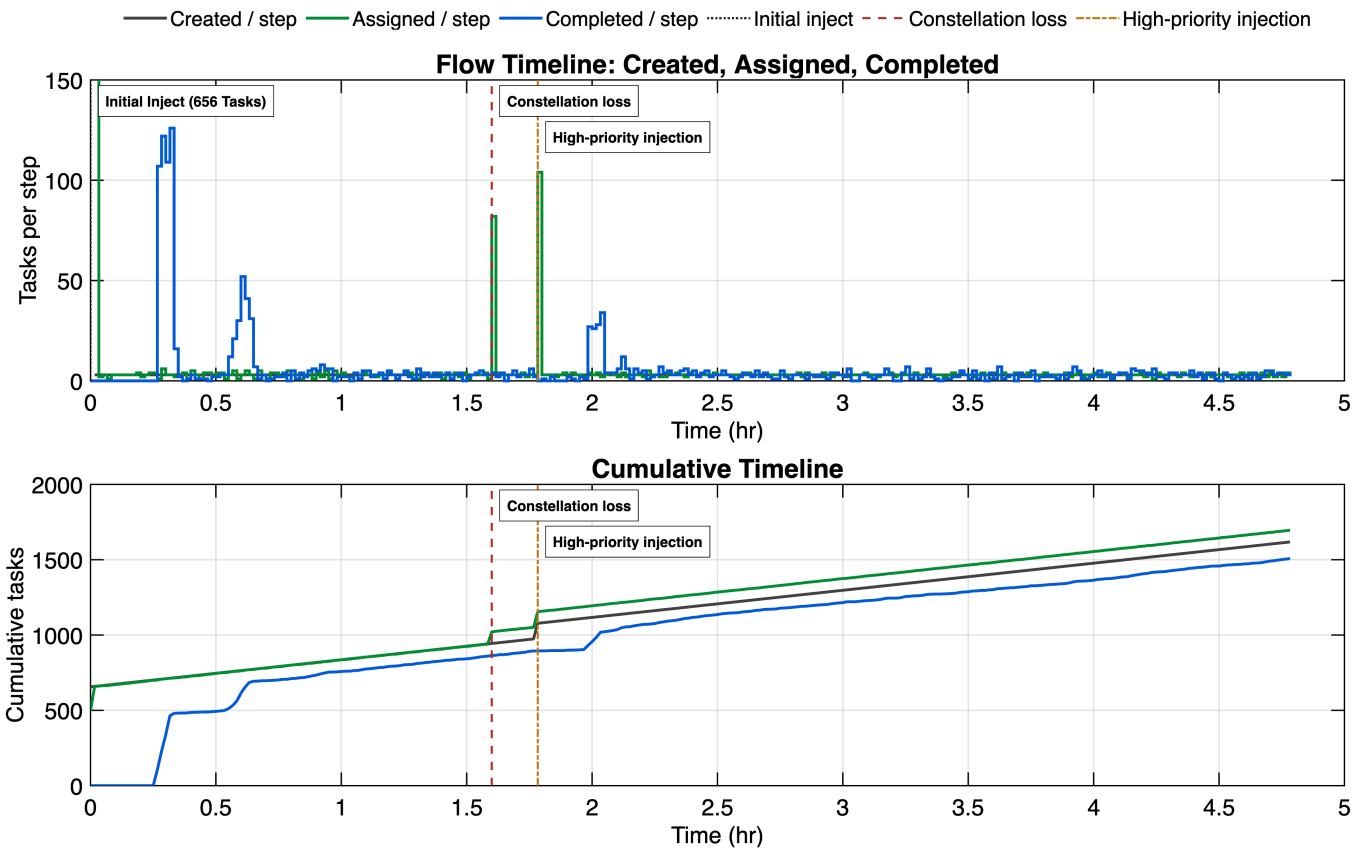

Figure 4: Created, assigned, and completed task flow and cumulative stock over the degraded SEUS run. Vertical event markers indicate the initial workload insertion, the step-96 constellation loss, and the step-107 high-priority injection.

marker is therefore the execution-time consequence of immediate re-prioritization, not delayed admission.

The aggregate mission accounting reinforces this interpretation. Across the full horizon, the framework completes 1,508 of 1,618 created tasks, yielding a 93.2% task completion rate. The priority satisfaction index is higher, at 96.3%, which means the surviving constellation preserves mission value better than it preserves raw task volume. At the final step, no tasks remain pending; the 110-task created-completed gap is already absorbed as 14 assigned tasks and 96 tasks in execution. The residual gap is therefore finite-horizon work in progress rather than unclaimed demand.

Table 1 reports the same result in compact form. The high-priority injection is fully admitted and fully completed, with an HPI on-time completion rate of 1.0000 and missed-deadline rate of 0.0000. The admission audit finds no unassigned HPI tasks, no never-feasible HPI tasks, and no serviceable-but-unselected HPI tasks. The re-prioritization claim is therefore stronger than a throughput claim: after severe loss, every injected urgent task is both selected and completed before deadline.

Replanning remains bounded even though the two disruption epochs produce sharp queue updates. Queue churn peaks at 0.9921 at the step-96 loss event and reaches 0.8739 at the step-107 injection, but the run mean is only 0.0997. Natural queue removals account for

Table 1: Primary mission and high-priority injection outcomes for the representative degraded SEUS run.

Metric	*Value*
Mission-level performance	
Created / completed tasks	1618 / 1508
Task completion rate	0.9320
Priority satisfaction index	0.9632
On-time completion rate / missed-deadline rate	0.9320 / 0.0000
Weighted on-time completion / missed-deadline rate	0.9632 / 0.0000
High-priority injection response	
Injected high-priority tasks	101
Assigned / completed high-priority tasks	101 / 101
Assignment latency mean / p50 / p95	0/0/0 steps
Completion latency mean / p50 / p95	13.60/14.00/15.00 steps
High-priority on-time completion / missed-deadline rate	1.0000 / 0.0000
Accepted commitments at injection	104
High-priority commitments at injection	101
Replanning stability	
Queue churn mean / peak	0.0997 / 0.9921
Final pending / assigned / in-execution	0 / 14 / 96

1,504 completed-task departures, while decision-driven removals sum to 78 and decision-driven reorders are zero. The controller therefore reacts strongly at the moments where the mission state actually changes, but it does not exhibit persistent queue thrashing. Taken together, the heatmap and timeline show adaptive high-priority re-prioritization rather than nominal-service preservation: the degraded constellation loses volume, shifts coordination pressure toward the Florida Atlantic corridor, and spends its remaining feasible capacity on the mission-critical workload.

5 CONCLUSION

The SEUS test campaign demonstrated priority-sensitive mission continuity for degraded heterogeneous multi-orbit satellite constellations using a hybrid MFGT–MPC distributed fallback autonomy framework. In the representative stress test, a 500-satellite, 6-ground-station architecture was reduced to 146 satellites and 2 ground stations before 101 high-priority hurricane-response tasks were injected across a concentrated Florida Atlantic corridor. The result was a direct test of whether a damaged constellation could recognize urgent demand, redirect scarce capacity, and continue executing the mission without centralized reassignment or phase-specific retuning.

The framework preserved mission value under severe loss. The heatmap-driven MFGT layer converted backlog, agent scarcity, priority concentration, and service memory into regional coordination pressure, while the MPC layer translated that pressure into feasible onboard commitments. At the injection epoch, 101 of 104 accepted commitments targeted the urgent workload. All 101 high-priority tasks were assigned immediately, completed before deadline, and incurred zero missed deadlines. Across the full run, the framework completed 1,508 of 1,618 tasks, achieving a 93.2% task completion rate and a 96.3% priority satisfaction index.

When constellation capacity collapses and ground coordination is degraded, mission continuity depends on the surviving assets' ability to make timely, priority-aware decisions from imperfect information. Heatmap-led MFGT coordination coupled with MPC-based local feasibility provides an executable mechanism for off-nominal stress tests that are becoming much more likely in today's space domain. The distributed autonomy framework does not merely preserve task volume; it preserves mission value by directing remaining capacity toward the work that matters most.

ACKNOWLEDGMENTS

The views expressed in this paper are those of the authors and do not reflect the official policy or position of the United States Space Force, the Department of War, or the U.S. Government.

The authors gratefully express their sincere appreciation for the full financial support through the U.S. Department of War, the U.S. Space Force, the Air Force Institute of Technology, and the Secretary of the Air Force's Science, Technology, Engineering, and Mathematics PhD Fellowship. Generative AI was used only for limited grammatical editorial checks during preparation of the manuscript.

REFERENCES

[1] 2025. *ESA'S Annual Space Environment Report.* Technical Report. European Space Agency, Darmstadt.

[2] 2025. Flying through the biggest solar storm ever recorded. https://www.esa.int/Space_Safety/Space_weather/Flying_through_the_biggest_solar_storm_ever_recorded

[3] 2025. What a Solar Superstorm Could Mean for the US. https://www.usgs.gov/news/featured-story/what-a-solar-superstorm-could-mean-us

[4] 2026. Amazon Leo. https://www.aboutamazon.com/what-we-do/devices-services/amazon-leo

[5] 2026. *Authorization and Order: Space Exploration Holdings, LLC Request for Deployment and Operating Authority for the SpaceX Gen2 NGSO Satellite System.* Technical Report. Federal Communications Commission, Washington, D.C. 1–34 pages. https://docs.fcc.gov/public/attachments/DA-26-36A1.pdf

[6] 2026. Starlink: Satellite Technology. https://starlink.com/technology?srsltid=AfmBOorqRrP7e_UOPGvyNAF0zRHhYaqN1dIRzdKQ4PXtxyxH8a0d6qUr

[7] Nof Abuzainab and Walid Saad. 2018. Dynamic Connectivity Game for Adversarial Internet of Battlefield Things Systems. *IEEE Internet of Things Journal* 5, 1 (2 2018), 378–390. https://doi.org/10.1109/JIOT.2017.2786546

[8] Carles Araguz, Elisenda Bou-Balust, and Eduard Alarcón. 2018. Applying autonomy to distributed satellite systems: Trends, challenges, and future prospects. *Systems Engineering* 21, 5 (9 2018), 401–416. https://doi.org/10.1002/sys.21428

[9] Mohamed Khalil Ben-Larbi, Kattia Flores Pozo, Mirue Choi, Tom Haylok, Benjamin Grzesik, Andreas Haas, Dominik Krupke, Harald Konstanski, Volker Schaus, Sándor P. Fekete, Christian Schurig, and Enrico Stoll. 2021. Towards the automated operations of large distributed satellite systems. Part 2: Classifications and tools. *Advances in Space Research* 67, 11 (6 2021), 3620–3637. https://doi.org/10.1016/j.asr.2020.08.018

[10] Mohamed Khalil Ben-Larbi, Kattia Flores Pozo, Tom Haylok, Mirue Choi, Benjamin Grzesik, Andreas Haas, Dominik Krupke, Harald Konstanski, Volker Schaus, Sándor P. Fekete, Christian Schurig, and Enrico Stoll. 2021. Towards the automated operations of large distributed satellite systems. Part 1: Review and paradigm shifts. *Advances in Space Research* 67, 11 (6 2021), 3598–3619. https://doi.org/10.1016/j.asr.2020.08.009

[11] Tom Butash, Peter Garland, and Barry Evans. 2021. Non-geostationary satellite orbit communications satellite constellations history. *International Journal of Satellite Communications and Networking* 39, 1 (1 2021), 1–5. https://doi.org/10.1002/sat.1375

[12] Furong Chai, Qi Zhang, Haipeng Yao, Xiangjun Xin, Ran Gao, and Mohsen Guizani. 2023. Joint Multi-Task Offloading and Resource Allocation for Mobile Edge Computing Systems in Satellite IoT. *IEEE Transactions on Vehicular Technology* 72, 6 (6 2023), 7783–7795. https://doi.org/10.1109/TVT.2023.3238771

[13] Giacomo Curzi, Dario Modenini, and Paolo Tortora. 2020. Large Constellations of Small Satellites: A Survey of Near Future Challenges and Missions. *Aerospace* 7, 9 (9 2020), 1–18. https://doi.org/10.3390/AEROSPACE7090133

[14] Inigo del Portillo, Bruce G. Cameron, and Edward F. Crawley. 2019. A technical comparison of three low earth orbit satellite constellation systems to provide global broadband. *Acta Astronautica* 159 (6 2019), 123–135. https://doi.org/10.1016/j.actaastro.2019.03.040

[15] Inigo del Portillo, Sydney I. Dolan, Bruce G. Cameron, and Edward F. Crawley. 2023. Architectural Decisions for Communications Satellite Constellations to Maintain Profitability While Serving Uncovered and Underserved Communities. *International Journal of Satellite Communications and Networking* 41, 1 (1 2023), 82–97. https://doi.org/10.1002/sat.1464

[16] Boualem Djehiche, Alain Tcheukam, and Hamidou Tembine. 2017. Mean-Field-Type Games in Engineering. *AIMS Electronic Engineering* 1, 1 (11 2017), 18–73. https://doi.org/10.3934/ms.2017.1.18

[17] J. Mikael Eklund, Jonathan Sprinkle, and S. Shankar Sastry. 2012. Switched and symmetric pursuit/evasion games using online model predictive control with application to autonomous aircraft. *IEEE Transactions on Control Systems Technology* 20, 3 (5 2012), 604–620. https://doi.org/10.1109/TCST.2011.2136435

[18] Karthik Elamvazhuthi and Spring Berman. 2019. Mean-Field Models in Swarm Robotics: A Survey. *Bioinspiration & Biomimetics* 15, 1 (11 2019), 1–15. https://doi.org/10.1088/1748-3190/ab49a4

[19] Huilong Fan, Chongxiang Sun, Jun Long, Ling Li, Yakun Huo, and Shangpeng Wang. 2025. Graph-Driven Resource Allocation Strategies in Satellite IoT: A Cooperative Game-Theoretic Approach. *IEEE Internet of Things Journal* 12, 4 (2025), 3463–3481. https://doi.org/10.1109/JIOT.2024.3407123

[20] Xiying Fan, Di Liu, Mengxuan Qiu, Yingqi Li, Jiahao Huo, Haojin Li, and Chen Sun. 2025. Dynamic Prioritized Data Transmission Through Intersatellite Cooperation in LEO Constellations. *IEEE Internet of Things Journal* 12, 2 (2025), 1922–1932. https://doi.org/10.1109/JIOT.2024.3464531

[21] Shaohan Feng, Xiao Lu, Sumei Sun, Ekram Hossain, Guiyi Wei, and Zhengwei Ni. 2024. Covert Communication in Large-Scale Multi-Tier LEO Satellite Networks. *IEEE Transactions on Mobile Computing* 23, 12 (2024), 11576–11587. https://doi.org/10.1109/TMC.2024.3396793

[22] Yufang Gao, Zhi Ji, Kanglian Zhao, Tomaso De Cola, and Wenfeng Li. 2024. Game-Based Computation Offloading and Power Allocation for LEO Constellation Networks in Distributed and Dynamic Environment. *IEEE Internet of Things Journal* 11, 4 (2 2024), 7040–7058. https://doi.org/10.1109/JIOT.2023.3314650

[23] Nathaniel G. Gordon, Nesrine Benchoubane, Gunes Karabulut Kurt, and Gregory Falco. 2025. On the Role of Communications for Space Domain Awareness. *Journal of Aerospace Information Systems* 22 (12 2025), 1–12. https://doi.org/10.2514/1.i011629

[24] Mohamad A. Hady, Siyi Hu, Mahardhika Pratama, Jimmy Cao, and Ryszard Kowalczyk. 2025. Multi-Agent Reinforcement Learning for Autonomous Multi-Satellite Earth Observation: A Realistic Case Study. (11 2025). http://arxiv.org/abs/2506.15207

[25] B Haluk and H Bozma. 2010. Multi-robot Navigation with Limited Communication - Deterministic vs Game-Theoretic Networks. In *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, Taipei, 1825–1830.

[26] Joshua Holder, Spencer Kraisler, and Mehran Mesbahi. 2025. Centralized and Distributed Strategies for Handover-Aware Task Allocation in Satellite Constellations. *Journal of Guidance, Control, and Dynamics* 48, 6 (6 2025), 1201–1210. https://doi.org/10.2514/1.G008363

[27] Ye Hu, Walid Saad, T Charles Clancy, Jeffrey H Reed, Kekatos Vassilis, and Craig A Woolsey. 2021. *Game Theory and Meta Learning for Optimization of Integrated Satellite-Drone-Terrestrial-Communication Systems*. Ph.D. Dissertation. Virginia Polytechnic Institute and State University, Blacksburg.

[28] Han Huang, Jiajia Yu, Jie Chen, and Rongjie Lai. 2023. Bridging mean-field games and normalizing flows with trajectory regularization. *J. Comput. Phys.* 487 (8 2023). https://doi.org/10.1016/j.jcp.2023.112155

[29] Kuang Huang, Xuan Di, Qiang Du, and Xi Chen. 2020. A Game-Theoretic Framework for Autonomous Vehicles Velocity Control: Bridging Microscopic Differential Games and Macroscopic Mean Field Games. (12 2020). https://doi.org/10.3934/dcdsb.2020131

[30] Minyi Huang, Roland Malhame, and Peter Caines. 2006. Large Population Stochastic Dynamic Games: Closed-Loop McKean-Vlasov Systems and the Nash Certainty Equivalence Principle. *Communications in Information and Systems* 6, 3 (2006), 221–252.

[31] Pengyan Huang, Guangchen Wang, Shujun Wang, and Hua Xiao. 2024. A Mean-Field Game for a Forward-Backward Stochastic System with Partial Observation and Common Noise. *IEEE/CAA Journal of Automatica Sinica* 11, 3 (3 2024), 746–759. https://doi.org/10.1109/JAS.2023.124047

[32] Daisuke Inoue, Yuji Ito, Takahito Kashiwabara, Norikazu Saito, and Hiroaki Yoshida. 2021. Model Predictive Mean Field Games for Controlling Multi-Agent Systems. *arXiv* (8 2021), 1–10. http://arxiv.org/abs/2004.07994

[33] Chunxiao Jiang and Zhen Li. 2020. Decreasing Big Data Application Latency in Satellite Link by Caching and Peer Selection. *IEEE Transactions on Network Science and Engineering* 7, 4 (10 2020), 2555–2565. https://doi.org/10.1109/TNSE.2020.2994638

[34] Julian Barreiro-Gomez, Tyrone E. Duncan, and Hamidou Tembine. 2019. Linear-Quadratic Mean-Field-Type Games-Based Stochastic Model Predictive Control: A Microgrid Energy Storage Application. In *2019 American Control Conference (ACC)*. AACC, Philadelphia, 3224–3229.

[35] Yuhan Kang, Yifei Zhu, Dan Wang, Zhu Han, and Tamer Basar. 2024. Joint Server Selection and Handover Design for Satellite-Based Federated Learning Using Mean-Field Evolutionary Approach. *IEEE Transactions on Network Science and Engineering* 11, 2 (3 2024), 1655–1667. https://doi.org/10.1109/TNSE.2023.3328776

[36] Oltjon Kodheli, Eva Lagunas, Nicola Maturo, Shree Krishna Sharma, Bhavani Shankar, Jesus Fabian Mendoza Montoya, Juan Carlos Merlano Duncan, Danilo Spano, Symeon Chatzinotas, Steven Kisseleff, Jorge Querol, Lei Lei, Thang X. Vu, and George Goussetis. 2021. Satellite Communications in the New Space Era: A Survey and Future Challenges. , 70–109 pages. https://doi.org/10.1109/COMST.2020.3028247

[37] Jean Michel Lasry and Pierre Louis Lions. 2007. Mean field games. *Japanese Journal of Mathematics* 2, 1 (2 2007), 229–260. https://doi.org/10.1007/s11537-007-0657-8

[38] Mathieu Laurière and Ludovic Tangpi. 2022. Convergence of large population games to mean field games with interaction through the controls. *arXiv* (3 2022). http://arxiv.org/abs/2004.08351

[39] Wei Li, Luliang Jia, Quan Chen, and Yingwu Chen. 2024. A Game Theory-Based Distributed Downlink Spectrum Sharing Method in Large-Scale Hybrid Satellite Constellations. *IEEE Transactions on Communications* 72, 8 (2024), 4620–4632. https://doi.org/10.1109/TCOMM.2024.3375813

[40] Xin Lin, Aijun Liu, Chen Han, Xiaohu Liang, Kegang Pan, and Zhixiang Gao. 2023. LEO Satellite and UAVs Assisted Mobile Edge Computing for Tactical Ad-Hoc Network: A Game Theory Approach. *IEEE Internet of Things Journal* 10, 23 (12 2023), 20560–20573. https://doi.org/10.1109/JIOT.2023.3299950

[41] Pierre Louis Lions and Jean Michel Lasry. 2007. Large investor trading impacts on volatility. *Annales de l'Institut Henri Poincare (C) Analyse Non Lineaire* 24, 2 (2007), 311–323. https://doi.org/10.1016/j.anihpc.2005.12.006

[42] Mark W. Maier, Frank W. Gallagher, Karen St. Germain, Richard Anthes, Cinzia Zuffada, Robert Menzies, Jeffrey Piepmeier, David Di Pietro, Monica M. Coakley, and Elena Adams. 2021. Architecting the future of weather satellites. *Bulletin of the American Meteorological Society* 102, 3 (2021), E589–E610. https://doi.org/10.1175/BAMS-D-19-0258.1

[43] Outer Space Institute. 2026. *CRASH Clock*. https://outerspaceinstitute.ca/crashclock/ CRASH Clock value listed as 3.0 days on 20 March 2026.

[44] Pierre Degond, Michael Herty, and Jian-Guo Liu. 2017. Meanfield Games and Model Predictive Control. *Communications in Mathematical Sciences* 15, 5 (2 2017), 1403–1422. https://intlpress.com/JDetail/180626499443621889

[45] Tomas Pippia, Valentin Preda, Samir Bennani, and Tamas Keviczky. 2022. Reconfiguration of a satellite constellation in circular formation orbit with decentralized model predictive control. *arXiv* (1 2022). http://arxiv.org/abs/2201.10399

[46] Yongtao Su, Yaoqi Liu, Yiqing Zhou, Jinhong Yuan, Huan Cao, and Jinglin Shi. 2019. Broadband LEO satellite communications: Architectures and key technologies. *IEEE Wireless Communications* 26, 2 (4 2019), 55–61. https://doi.org/10.1109/MWC.2019.1800299

[47] Kathiravan Thangavel, Roberto Sabatini, Alessandro Gardi, Kavindu Ranasinghe, Samuel Hilton, Pablo Servidia, and Dario Spiller. 2023. Artificial Intelligence for Trusted Autonomous Satellite Operations. *Progress in Aerospace Sciences* 144 (12 2023). https://doi.org/10.1016/j.paerosci.2023.100960

[48] Kathiravan Thangavel, Dario Spiller, Roberto Sabatini, Stefania Amici, Nicolas Longepe, Pablo Servidia, Pier Marzocca, Haytham Fayek, and Luigi Ansalone. 2023. Trusted Autonomous Operations of Distributed Satellite Systems Using Optical Sensors †. *Sensors* 23, 6 (3 2023). https://doi.org/10.3390/s23063344

[49] Dezhi Wang, Wei Wang, Yuhan Kang, and Zhu Han. 2022. Dynamic Data Offloading for Massive Users in Ultra-dense LEO Satellite Networks based on Stackelberg Mean Field Game. In *IEEE INFOCOM WKSHPS: PerAI-6G 2022: Pervasive Network Intelligence for 6G Networks*. Institute of Electrical and Electronics Engineers Inc. https://doi.org/10.1109/INFOCOMWKSHPS54753.2022.9798204

[50] Dezhi Wang, Wei Wang, Yuhan Kang, and Zhu Han. 2023. Distributed Data Offloading in Ultra-Dense LEO Satellite Networks: A Stackelberg Mean-Field Game Approach. *IEEE Journal of Selected Topics in Signal Processing* 17, 1 (2 2023), 112–127. https://doi.org/10.1109/JSTSP.2022.3226400

[51] Lixiang Wang, Dong Ye, Xianren Kong, Ming Liu, and Yan Xiao. 2024. Decentralized Receding Horizon Control for Satellite Cluster Reconfigurations With Successive Convexification Method. *IEEE Trans. Aerospace Electron. Systems* 60, 5 (2024), 5920–5936. https://doi.org/10.1109/TAES.2024.3398607

[52] Yao Wang, Chungang Yang, Tong Li, Xinru Mi, Lixin Li, and Zhu Han. 2024. A Survey on Mean-Field Game for Dynamic Management and Control in Space-Air-Ground Network. *IEEE Communications Surveys and Tutorials* 26, 4 (2024), 2798–2835. https://doi.org/10.1109/COMST.2024.3393369

[53] Yang Wu, Guyu Hu, Fenglin Jin, and Jiachen Zu. 2019. A Satellite Handover Strategy Based on the Potential Game in LEO Satellite Networks. *IEEE Access* 7 (2019), 133641–133652. https://doi.org/10.1109/ACCESS.2019.2941217

[54] Zhuochen Xie, Lu Ma, and Xuwen Liang. 2012. Unlicensed spectrum sharing game between LEO satellites and terrestrial cognitive radio networks. *Chinese Journal of Aeronautics* 25, 4 (8 2012), 605–614. https://doi.org/10.1016/S1000-9361(11)60425-1

[55] Jiapeng Xu, Xiang Chen, Ying Tan, and Guoxiang Gu. 2024. Robust Mean-Field Games With Partial Observations: A Complementary Strategy. *IEEE Trans. Automat. Control* 69, 12 (2024), 8766–8773. https://doi.org/10.1109/TAC.2024.3419002

[56] Yun Xu, Yulin Zhang, and Li Fan. 2024. Autonomous semi-major axis adjustment for mega constellation continuous coverage. *Advances in Space Research* 73, 11 (6 2024), 5582–5594. https://doi.org/10.1016/j.asr.2023.07.016

[57] Weiyi Yang, Yingwu Chen, Xiaolu Liu, Jun Wen, and Lei He. 2025. Distributed Satellites Dynamic Allocation for Grids with Time Windows: A Potential Game Approach. *Journal of Cleaner Production* (3 2025), 1–19. http://arxiv.org/abs/2503.08385

[58] Yafeng Zhan, Peng Wan, Chunxiao Jiang, Xiaohan Pan, Xi Chen, and Song Guo. 2020. Challenges and Solutions for the Satellite Tracking, Telemetry, and Command System. *IEEE Wireless Communications* 27, 6 (12 2020), 12–18. https://doi.org/10.1109/MWC.001.2000089

[59] Peixin Zhang, Wenting Wei, Kun Wang, Lizhe Liu, Liang Qin, and Celimuge Wu. 2024. Coalition game-based clustering algorithm for LEO satellite networks. In *GLOBECOM 2024 - 2024 IEEE Global Communications Conference*. IEEE Global Communications Conference, 5042–5047. https://doi.org/10.1109/GLOBECOM52923.2024.10901516

[60] Tao Zhang, Yiji Zhu, Dongying Ma, Chaoyong Li, and Xiaodong Wang. 2024. Toward Rapid and Optimal Strategy for Swarm Conflict: A Computational Game Approach. *IEEE Trans. Aerospace Electron. Systems* 60, 3 (6 2024), 3108–3120. https://doi.org/10.1109/TAES.2024.3361436

[61] Liancheng Zheng, Xuemei Wang, Feng Li, Zebing Mao, Zhen Tian, Yanhong Peng, Fujiang Yuan, and Chunhong Yuan. 2025. A Mean-Field-Game-Integrated MPC-QP Framework for Collision-Free Multi-Vehicle Control. *Drones* 9, 5 (5 2025). https://doi.org/10.3390/drones9050375

[62] Renjun Zheng, Haibo Wang, Matthieu De Mari, Miao Cui, Xiaoli Chu, and Tony Q.S. Quek. 2021. Dynamic Computation Offloading in Ultra-Dense Networks Based on Mean Field Games. *IEEE Transactions on Wireless Communications* 20, 10 (10 2021), 6551–6565. https://doi.org/10.1109/TWC.2021.3075028

[63] Zixuan Zheng, Jian Guo, and Eberhard Gill. 2018. Onboard mission allocation for multi-satellite system in limited communication environment. *Aerospace Science*

and Technology 79 (8 2018), 174–186. https://doi.org/10.1016/j.ast.2018.05.022

[64] Zejian Zhou and Hao Xu. 2022. A Novel Mean-Field-Game-Type Optimal Control for Very Large-Scale Multiagent Systems. *IEEE Transactions on Cybernetics* 52, 6 (6 2022), 5197–5208. https://doi.org/10.1109/TCYB.2020.3028267