# Large-Scale Continual Scheduling and Execution for Dynamic Distributed Satellite Constellation Observation Allocation

Extended Abstract

Itai Zilberstein
Carnegie Mellon University & Jet Propulsion Laboratory, California Institute of Technology
Pittsburgh, PA, USA
izilbers@cs.cmu.edu

Steve Chien
Jet Propulsion Laboratory, California Institute of Technology
Pasadena, CA, USA
steve.a.chien@nasa.jpl.gov

## ABSTRACT

The size and capabilities of Earth-observing satellite constellations are rapidly increasing. Leveraging distributed onboard control, we can enable novel time-sensitive measurements and responses. However, deploying autonomy to large multiagent satellite systems necessitates algorithms with efficient computation and communication. We tackle this challenge and propose new, online algorithms for large-scale *dynamic distributed constraint optimization problems (DDCOP)*. We present the *Dynamic Multi-Satellite Constellation Observation Scheduling Problem (DCOSP)*, a new formulation of DDCOPs that models integrated scheduling and execution. We construct an omniscient offline algorithm to compute the novel optimality condition of DCOSP and present the *Dynamic Incremental Neighborhood Stochastic Search (D-NSS)* algorithm, an incomplete online decomposition-based DDCOP approach. We show through simulation that D-NSS converges to near-optimal solutions and outperforms DDCOP baselines in terms of solution quality, computation time, and message volume. Our work forms the foundation of the largest in-space demonstration of distributed multiagent AI to date: the NASA FAME mission.

## KEYWORDS

Distributed Constraint Optimization; Scheduling; Satellite Operations

**ACM Reference Format:**
Itai Zilberstein and Steve Chien. 2026. Large-Scale Continual Scheduling and Execution for Dynamic Distributed Satellite Constellation Observation Allocation: Extended Abstract. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 3 pages. https://doi.org/10.65109/JCYH5778

## 1 INTRODUCTION

There has been a proliferation of Earth-observing spacecraft in recent years and advancements in their capabilities to act as autonomous agents [1, 2, 4, 5, 8, 12]. Large observation systems result in shorter revisit times to the target observation locations; this is crucial for rapid responses to dynamic events such as natural disasters.



*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.* © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). https://doi.org/10.65109/JCYH5778

This work focuses on distributed online scheduling that can efficiently coordinate the actions of hundreds or thousands of spacecraft continually in a dynamic environment. We model a large-scale satellite constellation as a multiagent system and focus on optimizing observation completion during an overlapping scheduling and execution horizon. We present the *Dynamic Multi-Satellite Constellation Observation Scheduling Problem (DCOSP)*, a *dynamic distributed constraint optimization problem (DDCOP)* formulation of satellite scheduling. DCOSP extends the basic DDCOP model: it has a novel optimality condition that takes into account integrated scheduling and execution. In addition, realistic problem instances consist of millions of continuously changing variables. Agents must react to dynamics while expending limited time, memory, and communication. We extend the *Neighborhood Stochastic Search (NSS)* algorithm [13] to a dynamic variation referred to as *Dynamic Incremental Neighborhood Stochastic Search (D-NSS)*. Our work will be leveraged in the largest in-space demonstration of multiagent AI to date, beginning in 2026. The NASA FAME demonstration involves over 60 participating spacecraft that will dynamically coordinate to observe Earth phenomena [3]. The contributions of this work are

(1) formulating the real-world application of dynamic satellite scheduling as a DDCOP with a unique optimality condition that models task execution,
(2) constructing an omniscient, offline optimal solver for the dynamic satellite scheduling problem, and
(3) presenting the *Dynamic Incremental Neighborhood Stochastic Search algorithm*, a scalable incomplete DDCOP approach.

## 2 PROBLEM DEFINITION

The *Multi-Satellite Constellation Observation Scheduling Problem (COSP)* [13] consists of a set of agents (satellites) $A$ tasked with satisfying observation requests $R$ for ground targets within a scheduling horizon $H$. Each agent $a \in A$ manages a set of potential tasks $S_a$ and Boolean decision variables $x \in \mathcal{X}_a$. Scheduling is subject to strict resource constraints: agents cannot overlap tasks or downlinks, must adhere to memory capacities, and must downlink data. Unlike standard DCOPs, the COSP constraint graph is dense, with degrees of $\Omega(|A| \cdot |R|)$, and only locally known. Agents are oblivious to the variables and constraints of their peers.

We extend COSP to *Dynamic COSP (DCOSP)*, modeled as a sequence of COSP instances $\delta = \{\delta_t\}_{t=0}^{T}$. Since the problem is online and reactive, traditional DDCOP utility, which sums the utility of individual static instances, is insufficient. In DCOSP, utility is defined by the number of requests executed instead of scheduled ones. We define the proposition executed$(x)$ to track successful completion:

$$\text{executed}(x) = x \cdot \mathbb{I}[\exists_t\ x = x(s) \wedge h(s) \cap \bar{h}(\delta_t) \neq \emptyset]$$

where $\bar{h}(\delta_t)$ is the interval during which the problem remains static. The total utility $\mathcal{F}(\mathcal{X}^\delta)$ measures the executed requests:

$$\mathcal{F}(\mathcal{X}^\delta) = \sum_{r \in R^\delta} \left[ 1 - \prod_{x \in \mathcal{X}_r^\delta} (1 - \text{executed}(x)) \right]$$

The dynamics render many existing DCOP methods computationally intractable due to high graph dimensionality and the requirement that agents act reactively without global knowledge of neighboring variables.

## 3 ALGORITHMS

### 3.1 Optimality and Omniscient Benchmarking

Finding an optimal solution to DCOSP is challenging because agents lack prior knowledge of problem dynamics. However, we can establish a performance upper bound by collapsing a DCOSP $\delta$ into a static DCOP $\delta'$ using *omniscient* knowledge. By restricting the variables across $\delta$ to the tasks that persist across the horizon, we ensure that $\mathcal{F}(\mathcal{X}^\delta) = \mathcal{F}(\mathcal{X}^{\delta'})$. While the omniscient solver is not deployable, it serves as a benchmark for evaluating practical online algorithms.

### 3.2 Dynamic Incremental NSS (D-NSS)

We propose *Dynamic Incremental Neighborhood Stochastic Search* (D-NSS), an efficient, local-search algorithm designed for the scale and resource constraints of satellite constellations. D-NSS operates in three phases: (i) partitions the constraint graph using a decomposition heuristic, (ii) repairs schedules by removing obsolete tasks and inserting new ones when dynamics occur, and (iii) iteratively refines solutions within a partition. D-NSS is highly scalable, with a per-iteration computation and communication complexity of $O(|A_\mathcal{N}| \cdot |R_\mathcal{N}|)$ where $|A_\mathcal{N}|$ and $|R_\mathcal{N}|$ are the maximum number of agents and requests in a partition.

## 4 EXPERIMENTS

### 4.1 Experimental Setup

We evaluate D-NSS with the *Geometric Neighborhood Decomposition Heuristic* [13] against several baselines: 0-NSS (recomputing from scratch), D-DSA (*Distributed Stochastic Search Algorithm* with repair) [6, 10], 0-DSA, and non-communicating *greedy* and *random* solvers. We model two satellite constellations: *Planet* (200 satellites) and *Walker* (108 satellites), with ground station downlinks and a 24-hour scheduling horizon.

### 4.2 Results

For small instances, we benchmark against the omniscient optimal solution obtained via centralized branch-and-bound. As shown in Table 1, D-NSS achieves the lowest optimality gap while utilizing an order of magnitude less computation time and communication than DSA baselines.

In large-scale scenarios with thousands of requests, D-NSS consistently outperforms baselines in satisfaction rate while maintaining significantly lower computation. D-NSS reduces message volume by up to two orders of magnitude compared to D-DSA due to its neighborhood-based decomposition. Convergence analysis also reveals that D-NSS is more stable than from-scratch solvers. The repair function allows D-NSS to maintain high satisfaction even immediately following a problem change.

| Algorithm | Optimality Gap (%) | Time (ms) | Messages (KB) |
|---|---|---|---|
| *Walker Constellation* | | | |
| Random | 14.95 | < 1 | 0 |
| Greedy | 15.60 | < 1 | 0 |
| **D-NSS** | **0.14** | **5.2** | **240.2** |
| 0-NSS | 0.48 | 6.1 | 400.0 |
| D-DSA | 1.17 | 58.2 | 10,459.5 |
| 0-DSA | 1.22 | 54.0 | 7,789.0 |
| *Planet Constellation* | | | |
| Random | 2.53 | < 1 | 0 |
| Greedy | 8.37 | < 1 | 0 |
| **D-NSS** | **1.87** | **<1** | **7.3** |
| 0-NSS | 2.59 | < 1 | 13.2 |
| D-DSA | 4.22 | 5.3 | 980.6 |
| 0-DSA | 3.80 | 4.7 | 718.4 |

**Table 1: Algorithm performance on dynamic scenarios with up to 1000 requests for the Walker and Planet constellations.**

## 5 CONCLUSION

Scaling dynamic DCOP algorithms to large-scale problems like DCOSP remains a significant challenge. We present D-NSS, a decomposition and repair-based algorithm that achieves an order of magnitude reduction in message volume and runtime compared to baselines while achieving near-optimal utility in dynamic environments. The D-NSS algorithm is generalizable to other domains such as multiagent pathfinding [9] and mobile sensor teams [7]. The approach will be demonstrated in 2026 through NASA's FAME mission [3, 11]. By enabling distributed onboard coordination, this work lays the foundation for more responsive Earth-observing systems capable of managing time-critical events like natural disasters.


## ACKNOWLEDGMENTS

Portions of the research were carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004). Government sponsorship acknowledged. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program (DGE2140739). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.


## REFERENCES


[1] Steve Chien, Alberto Candela, Itai Zilberstein, David Rijlaarsdam, Tom Hendrix, and Aubrey Dunne. 2024. Leveraging commerical assets, edge computing, and near real-time communications for an enhanced New Observing Strategies (NOS) flight demonstration. In *Proceedings of the IEEE Geoscience and Remote Sensing Symposium*.

[2] Steve Chien, Rob Sherwood, Daniel Tran, Benjamin Cichy, Gregg Rabideau, Rebecca Castano, Ashley Davis, Dan Mandl, Stuart Frye, Bruce Trout, et al. 2005.

Using autonomy flight software to improve science return on Earth Observing One. *Journal of Aerospace Computing, Information, and Communication* 2, 4 (2005), 196–216.

[3] Steve Chien, Itai Zilberstein, Alberto Candela, Domenico Barretta, David Rijlaarsdam, Tom Hendrix, Aubrey Dunne, Oriol Cortés Grauc, Alexandre Gol i Mestrec, Manel Pedra Bovec, Oriol Aragon, Juan Puig Miquel, Arvind Subramanian, Vishesh Vatsal, Adithya Kothandhapani, Jad Mogannam, and Mitchell Scher. 2025. Multi-asset New Observing Systems flight demonstration. In *Proceedings of the International Conference on Space Operations*.

[4] Steve Chien, Itai Zilberstein, Alberto Candela, David Rijlaarsdam, Amaury Perrocheau, Aubrey Dunne, Tom Hendrix, Oriol Cortés Grauc, Alexandre Gol i Mestrec, Manel Pedra Bovec, Oriol Aragon, and Juan Puig Miquel. 2025. Flight of dynamic targeting on CogniSAT-6 - Update. In *Proceedings of the International Conference on Space Operations*.

[5] NewSpace. 2025. NewSpace Constellations. https://www.newspace.im. Accessed: 2025-05-01.

[6] Shreya Parjan and Steve A. Chien. 2023. Decentralized observation allocation for a large-scale constellation. *Journal of Aerospace Information Systems* (2023), 1–15.

[7] Arseniy Pertzovsky, Roie Zivan, and Noa Agmon. 2024. Collision avoiding Max-Sum for mobile sensor teams. *Journal of Artificial Intelligence Research* 79 (2024), 1281–1311.

[8] Gregg Rabideau, Joseph Russino, Andrew Branch, Nihal Dhamani, Tiago Stegun Vaquero, Steve Chien, Jean-Pierre de la Croix, and Federico Rossi. 2025. Planning, scheduling, and execution on the Moon: the CADRE technology demonstration mission. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems*.

[9] Oren Salzman and Roni Stern. 2020. Research challenges and opportunities in multi-agent path finding and multi-agent pickup and delivery problems. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems*. 1711–1715.

[10] Weixiong Zhang, Guandong Wang, Zhao Xing, and Lars Wittenburg. 2005. Distributed stochastic search and distributed breakout: Properties, comparison and applications to constraint optimization problems in sensor networks. *Artificial Intelligence* 161, 1-2 (2005), 55–87.

[11] Itai Zilberstein, Alberto Candela, and Steve Chien. 2025. Federated autonomous operations: A New paradigm for large-scale observation systems. In *Proceedings of the International Conference on Space Operations*.

[12] Itai Zilberstein, Alberto Candela, Steve Chien, David Rijlaarsdam, Tom Hendrix, Léonie Buckley, and Aubrey Dunne. 2024. Demonstrating onboard inference for Earth science applications with spectral analysis algorithms and deep learning. In *Proceedings of the International Symposium on Artificial Intelligence, Robotics and Automation in Space*.

[13] Itai Zilberstein, Ananya Rao, Matthew Salis, and Steve Chien. 2025. Decentralized, decomposition-based observation scheduling for a large-scale satellite constellation. *Journal of Artificial Intelligence Research* 82 (2025), 169–208.