

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

Itai Zilberstein

Carnegie Mellon University
Jet Propulsion Laboratory, California Institute of
Technology
izilbers@cs.cmu.edu

Steve Chien

Jet Propulsion Laboratory, California Institute of
Technology
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 satellites requires efficient computation and communication. This work tackles the challenge of efficiently scheduling observations for hundreds of satellites in a dynamic, large-scale problem with millions of variables. We present the *Dynamic Multi-Satellite Constellation Observation Scheduling Problem* (DCOSP), a new formulation of Dynamic Distributed Constraint Optimization Problems (DDCOP) that models integrated scheduling and execution. DCOSP has a novel optimality condition for which we construct an omniscient offline algorithm for its computation. We also present the *Dynamic Incremental Neighborhood Stochastic Search algorithm* (D-NSS), an incomplete online decomposition-based DDCOP algorithm that repairs and solves sub-problems when problem dynamics occur. 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. As part of the NASA FAME mission, DCOSP and D-NSS will be the foundation of the largest in-space demonstration of distributed multi-agent AI to date.

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. In *Appears at the International Workshop on Autonomous Agents and Multi-Agent Systems for Space Applications (MASSpace-26)*. Held as part of the Workshops at the 25th International Conference on Autonomous Agents and Multiagent Systems., Paphos, Cyprus, May 2026, IFAAMAS, 10 pages.

1 INTRODUCTION

There has been a proliferation of Earth-observing spacecraft in recent years, including advancements in their capabilities to act as autonomous agents. Reduced launch costs have led to constellations composed of hundreds or thousands of spacecraft that can monitor

Earth phenomena [31]. Large observation systems result in shorter revisit times to the target observation locations. Reduced revisit times are crucial for rapid responses to dynamic events such as natural disasters. New spacecraft also possess edge hardware capable of performing more intensive computation onboard, including neural network execution and even planning [7, 8, 10, 40, 51]. The advancement of inter-satellite links (ISL) has enabled persistent communications between spacecraft and stations on the ground.

These capabilities support observation campaigns that require time-sensitive and coordinated measurements. An example is global monitoring of all volcanic activity or flooding events. However, without consistent observation and the ability to react to dynamic events, key measurements of these processes may be missed. For example, wildfire monitoring requires 30-minute updates to be useful to ground responders [23]. Centralized scheduling on the ground suffers from latency and may not be able to meet these timing constraints. Therefore, to successfully tackle these campaigns, we require large-scale observation systems with reduced revisit times and distributed autonomy so that agents can operate online without ground control.

This work focuses on distributed online scheduling that can efficiently coordinate the actions of hundreds or thousands of spacecraft continually as problem dynamics change. Coordinating the actions of a large-scale constellation requires reasoning about agents with varying capabilities, constraints, and visibility of Earth targets while managing limited computational resources. Satellites share CPU and RAM with other critical flight software. Large volumes of communication in space are also unreliable and may even carry a financial cost [25]. These limitations make solving a static problem challenging. However, any operational solution must solve dynamic problems that change over time, which further increases the complexity. In these scenarios, the satellite constellation is fixed, yet the observation requests change. These changing observation requests alter the constraints of the problem. We desire continual scheduling solutions that are lightweight yet can effectively handle these problem dynamics.

Operational satellite observation scheduling uses centralized paradigms [43], and many research efforts have concentrated on centralized solutions [3, 30]. Centralized approaches may be insufficient for dynamic situations that require real-time response due to latencies to the ground. Centralized approaches are also vulnerable to single-point failures that would result in a non-operational constellation.

We model a large-scale satellite constellation as a multi-agent system (MAS) and focus on continual decentralized scheduling that optimizes observation completion during an overlapping scheduling and execution horizon. We present the *Dynamic Multi-Satellite*

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

Appears at the International Workshop on Autonomous Agents and Multi-Agent Systems for Space Applications (MASSpace-26). Held as part of the Workshops at the 25th International Conference on Autonomous Agents and Multiagent Systems., S. Chien, G. Picard, I. Zilberstein (Chairs), May 2026, Paphos, Cyprus. © 2026 Copyright held by the owner/author(s).

Constellation Observation Scheduling Problem (DCOSP), which is a dynamic distributed constraint optimization problem (DDCOP). DCOSP is a novel application of a DDCOP for dynamic satellite observation scheduling, and extends DDCOPs in multiple ways. The assumptions of DCOSP differ from those of a DDCOP to reflect real-world constraints. We assume agents are aware of the existence of other agents, but do not have access to detailed capabilities or state information. This assumption allows our approach to apply to scenarios with intermittent connectivity, limited bandwidth, or security-mandated communication restrictions. We also formulate a novel optimality condition for a DDCOP that takes into account integrated scheduling and execution. In addition, DDCOP solutions tend to rely heavily on computation and communication, especially when problem dynamics are volatile, which prevents application to DCOSP. These challenges make current DDCOP solvers insufficient for DCOSP.

DCOSP problem instances are much larger than typical DDCOP problems examined in previous literature. DCOSP instances consist of millions of variables that change over time. Agents must react to problem dynamics while expending limited time, memory, and communication. We extend the *Neighborhood Stochastic Search* algorithm (NSS) to a dynamic variation referred to as *Dynamic Incremental Neighborhood Stochastic Search* (D-NSS). NSS decomposes the global problem into smaller sub-problems, producing global solutions more efficiently. D-NSS uses the same foundation and leverages repairing previous solutions to efficiently handle problem dynamics.

The contributions of this work are

- (1) a novel formulation of 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 solution to the dynamic satellite scheduling problem, and
- (3) presenting the *Dynamic Incremental Neighborhood Stochastic Search algorithm*, a scalable incomplete DDCOP approach.

We evaluate the effectiveness of the approaches on large-scale real-world scenarios as well as analyze the challenges of deploying existing DDCOP solutions to this problem. DCOSP and D-NSS will be leveraged in the largest in-space demonstration of multi-agent AI to date, beginning in 2026. The NASA FAME demonstration involves over 60 participating spacecraft that will dynamically coordinate to observe Earth phenomena [9].

A version of this work appears as an extended abstract in the proceedings of AAMAS 2026 [52].

2 RELATED WORK

2.1 Satellite Observation Scheduling

Satellite observation scheduling is typically framed as an optimization problem that involves geometric reasoning, downlink scheduling, and constraint-based task allocation. The majority of research efforts and operational work have focused on centralized solutions to satellite observation scheduling [1, 3, 6, 11, 17, 18, 30, 43–45, 47]. There is limited work on decentralized scheduling approaches, and these mainly focus on static problems. Examples include auction-based methods [36, 37] and heuristic search-based methods [4, 5, 33, 53]. The work of Zilberstein, Rao, Salis, and Chien

proposed the *Multi-Satellite Constellation Observation Scheduling Problem* (COSP), a DCOP formulation of observation allocation [53]. We extend this work by formulating the *Dynamic Multi-Satellite Constellation Observation Scheduling Problem* (DCOSP), which is a novel DDCOP formulation of the problem.

2.2 Dynamic Distributed Constraint Optimization

Distributed constraint optimization problems (DCOP) have modeled applications including mobile sensor teams [34], smart grids [14], and satellite scheduling [53]. Solutions to distributed constraint optimization problems tend to be intensive in computation and communication, making deployment to agents with limited computation challenging. Optimal solutions have exponential complexities [16, 19, 28, 29, 35]. Incomplete DCOP algorithms are more efficient, yet typically rely on agents communicating with all neighboring agents in the constraint graph, resulting in large complexities when constraint graphs are fully connected [26, 32, 46, 49]. The *Neighborhood Stochastic Search* (NSS) algorithm, which iteratively improves sub-problem solutions, has been shown to solve large-scale distributed satellite observation scheduling with limited computation and communication [53]. We extend the NSS algorithm to the *dynamic* DCOP setting (DDCOP) to perform scalable and effective dynamic observation scheduling.

Dynamic distributed constraint optimization problems (DDCOP) [24] extend DCOPs to capture problem changes. A standard DDCOP is composed of a sequence of T static DCOPs where an optimal solution is obtained by solving each of the T DCOPs optimally. DDCOP solutions are inherently online algorithms as a system reacts to changes. Most work has adapted common DCOP algorithms to dynamic variations that inherit computational complexities [2, 12, 22, 27, 41, 48, 54].

When it comes to the application of DDCOPs to satellite operations, there are limitations with the standard definition. When solving static DCOPs, it is possible to assume that solutions are found prior to the execution horizon. However, DDCOPs cannot always make this assumption as dynamics may occur during the execution horizon. This is true for satellite operations; utility is obtained by taking an observation, not scheduling one. Therefore, when the problem changes so that a requested task is removed from the problem, having it scheduled prior to execution may not be optimal. In addition, due to consumptive resources, the starting state of subsequent DCOPs is driven by previous solutions.

3 PROBLEM DEFINITION

3.1 Multi-Satellite Constellation Observation Scheduling

We outline the *Multi-Satellite Constellation Observation Scheduling Problem* (COSP) and discuss the related challenges. COSP is defined by the following sets.

- $H = [h_s, h_e]$: the scheduling horizon.
- A : the set of agents in which each agent is a satellite in the constellation.
- \mathcal{T} : the set of point targets on Earth defined by a latitude and longitude.

- R : the set of requests where each request is defined by the target to observe, $\tau \in \mathcal{T}$, and when in the scheduling horizon to observe, $h \subset H$. Note that we index elements of a request r , such as the horizon, with the notation $h(r)$ and use this notation consistently for other variables.
- $S = \bigcup_{a \in A} S_a$: the set of tasks (also referred to as observations) where each S_a corresponds to the tasks of agent a . A task $s \in S_a$ is defined by the request being satisfied, $r \in R$, the interval required to schedule the task, $h \subset h(r)$, and the data volume required to take the observation, $m \in \mathbb{R}^+$.
- $\mathcal{X} = \bigcup_{a \in A} \mathcal{X}_a$: the set of Boolean decision variables where each \mathcal{X}_a corresponds to the variables of agent a . For each $s \in S_a$ we define the Boolean decision variable $x \in \mathcal{X}_a$ where $x = 1$ iff agent a schedules task s .
- $D = \bigcup_{a \in A} D_a$: the set of downlinks where each D_a corresponds to the downlinks of agent a . A downlink is defined by the maximum data volume downlinked, $m \in \mathbb{R}^+$, and the time interval of the downlink, $h \subset H$. We assume that all downlinks are mandatory.
- $C = \bigcup_{a \in A} C_a$: the set of constraints for each agent. Each agent is constrained by processing and data volume. An agent cannot execute two tasks at once and tasks cannot overlap with downlinks. An agent must also never exceed its memory capacity and all observations acquired must be downlinked at the earliest opportunity. Formally,

$$C_a = C_{D_a} \cup C_{S_a}.$$

We define

$$C_{D_a} = \bigcup_{d \in D_a} c_d$$

where

$$c_d = \sum_{s \in S_a^d} x(s) \cdot m(s) \leq \min(m(a), m(d)).$$

The set S_a^d contains the possible tasks for which the soonest downlink window in the future is d . The value $m(a)$ denotes the memory capacity of agent a . We define

$$C_{S_a} = \bigcup_{s, s' \in S_a} c_{s, s'}$$

where

$$c_{s, s'} = [x(s) \cdot x(s') + \mathbb{I}(h(s) \cap h(s') \neq \emptyset) \leq 1].$$

The objective of COSP is to maximize the number of requests satisfied subject to the constraints. A request is satisfied if a single observation for that request is completed. An optimal assignment of variables \mathcal{X}^* is defined as

$$\mathcal{X}^* = \arg \max_{\mathcal{X}} \mathcal{F}(\mathcal{X})$$

where

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

Here, \mathcal{X}_r is the set of variables such that $x = x(s)$ and $r(s) = r$.

COSP has been shown to be a challenging problem for DCOP methods. Typical COSP instances have hundreds of agents and thousands of requests, leading to millions of decision variables. In

addition, the constraint graph of COSP has high degrees on the order of $\Omega(|A| \cdot |R|)$. The constraint graph is also assumed to be only locally known to an agent. An agent a is oblivious to all tasks $s \notin S_a$ and variables $x \notin \mathcal{X}_a$. This is not consistent with standard DCOPs in which agents know all neighboring variables/agents in the constraint graph [13]. In COSP, it is assumed that agents know the existence of all other agents but have no knowledge of the variables of other agents. These factors make existing DCOP approaches that rely on agents communicating with neighboring agents in the constraint graph both computationally challenging due to the high degrees in the graph and inapplicable since we cannot assume agents know which agents they share constraints with.

3.2 Dynamic Multi-Satellite Constellation Observation Scheduling

We now present the *Dynamic Multi-Satellite Constellation Observation Scheduling Problem* (DCOSP) and discuss how previous challenges from COSP transfer to DCOSP.

A DDCOP consists of a set of T sequential DCOPs. We define DCOSP similarly as a set of COSP instances. Let δ_t be the COSP instance at time t . We then define the DCOSP, δ , as $\delta = \{\delta_t\}_{t=0}^T$. We assume that the agents have no prior knowledge of when or how the problem might change and must act reactively. We refer to the requests and variables of δ_t as R^{δ_t} and \mathcal{X}^{δ_t} . Note that δ_i depends on δ_j for $j < i$ since these prior DCOPs will determine the starting state of δ_i . For example, resource expenditure affects both current and future solutions. We assume that there is a globally known horizon for a DCOSP instance, $h(\delta) = [h_s(\delta), h_e(\delta)]$ and that the horizon of each COSP instance is $h(\delta_t) = [h_s(\delta_t), h_e(\delta_t)]$ where $h_s(\delta_t) \in h(\delta)$.

The utility of DCOSP is not the same as a typical DDCOP. In a typical DDCOP, the utility is defined as the sum of the utility functions of the individual DCOPs, which means that an optimal solution is obtained by solving each DCOP optimally in sequence. However, this formulation does not adequately capture DCOSP utility. We define the utility of DCOSP to be the number of requests that are satisfied, where satisfaction is determined by an observation for a request being *executed*. This utility rewards completing a task rather than just scheduling one. Due to the online nature of DCOSP, the scheduling horizon overlaps the execution horizon. Therefore, scheduling an observation does not guarantee that it will be executed. Consider that a task for request r is scheduled at time t_0 in δ_{t_0} to be executed at time t_i . If the task is then unscheduled at some δ_{t_j} where $0 < j < i$ then r will not be satisfied despite having a task scheduled in COSP instance δ_{t_0} .

To formally define this utility, we provide some useful definitions. Let $\bar{h}(\delta_t)$ be the unknown execution horizon of δ_t . We visually show $\bar{h}(\delta_t)$ in Figure 1. This is the horizon for which the problem is static and is defined by δ_t . Formally,

$$\bar{h}(\delta_t) = \begin{cases} [h_s(\delta_t), h_s(\delta_{t+1})] & \text{if } t < T \\ [h_s(\delta_t), h_e(\delta_t)] & \text{else (i.e. } t = T). \end{cases}$$

We then define the proposition executed(x) as

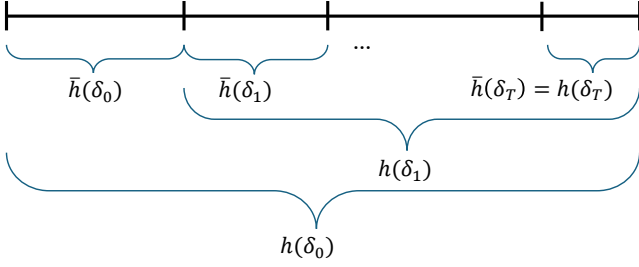


Figure 1: We illustrate $\bar{h}(\delta_t)$; the line spans the entire horizon of the DCOSP δ where $\bar{h}(\delta_t)$ shows the unknown interval the problem is static as defined by δ_t .

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

The above is equal to one if and only if $x = 1$, meaning task s was scheduled, and the horizon of s occurred during the time when the problem is static as defined by δ_t . This means that the task was scheduled when it was executed. Using this proposition, we can define the utility of DCOSP for an assignment of variables \mathcal{X}^δ over all time steps.

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

where $R^\delta = \bigcup_{t=0}^T R^{\delta_t}$ and $\mathcal{X}_r^\delta = \bigcup_{t=0}^T \mathcal{X}_r^{\delta_t}$. R^δ is the set of all requests in DCOSP δ and \mathcal{X}_r^δ are all tasks for request r over δ . This utility function rewards observations that are completed, not just scheduled.

The same challenges of solving COSP are multiplied when solving DCOSP. Solving a single instance of COSP is computationally challenging for most DCOP methods. Solutions that are linear in the maximum degree of the constraint graph suffer due to COSP instances having degrees $\Omega(|A| \cdot |R|)$. We again assume that the constraint graph is only partially known to an agent. Therefore, solutions to DCOSP need to conform to this assumption that cross-agent edges in the constraint graph are unknown and to be efficient in computation and communication.

4 ALGORITHMS

4.1 Obtaining an Optimal Solution to DCOSP

There is no clear online algorithm for computing an optimal DCOSP solution. This is due to agents having no prior knowledge of problem dynamics and optimally solving each individual COSP instance not necessarily constituting an optimal DCOSP solution. However, we can obtain an optimal solution by collapsing DCOSP into a single DCOP. This relies on using omniscient knowledge of the problem dynamics and is therefore not feasible in actuality.

We can collapse a DCOSP δ into a DCOP δ' to reason only about the observations that matter. The following are the key set constructions of δ' :

$$S^{\delta'} = \{s \in S^\delta \mid \exists t h(s) \cap \bar{h}(\delta_t) \neq \emptyset\} \text{ and}$$

$$\mathcal{X}^{\delta'} = \{x \in \mathcal{X}^\delta \mid \exists t x = x(s) \wedge h(s) \cap \bar{h}(\delta_t) \neq \emptyset\}.$$

By construction, $\text{executed}(x) = x$ if and only if $x \in \mathcal{X}^{\delta'}$ and $\text{executed}(x) = 0$ otherwise. Therefore, we obtain the following equivalence.

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

Solving δ' optimally results in an optimal solution to the DCOSP δ . By definition, δ' is a static COSP instance. Therefore, we can employ any complete algorithm to obtain an optimal solution to δ by solving δ' . Although the omniscient solver is not deployable, it serves as a key performance upper bound for benchmarking practical, online algorithms such as D-NSS. However, even solving a single DCOP is NP-Hard [29]. Therefore, for large DCOSP instances, we use an incomplete solver to obtain a near-optimal solution since any exponential time search is infeasible.

4.2 Dynamic Incremental Neighborhood Stochastic Search

In this section, we present the *Dynamic Incremental Neighborhood Stochastic Search* algorithm (D-NSS) shown in Algorithm 1. Due to the scale of DCOSP and the computational constraints of satellites, we require DDOP algorithms that are both efficient and can reason about the dynamic nature of the problem. D-NSS extends NSS to address the drawbacks of general DDOP algorithms. NSS is an iterative algorithm where at each iteration agents stochastically update their variable assignments based on the assignments of agents they communicate with. This iterative procedure is shared with other algorithms such as DSA. However, NSS relies on a decomposition heuristic $\Upsilon : A \times S \rightarrow \{0, 1\}$ to generate a subproblem \mathcal{N} for neighborhoods of agents to solve. This subproblem is a smaller DCOP consisting of requests $R_{\mathcal{N}}$ and agents $A_{\mathcal{N}}$. D-NSS continually computes subproblems, repairs local solutions between problem instances, and reasons about prior scheduling and execution in the search and repair phases.

We leverage the Geometric Neighborhood Decomposition heuristic (GND) to create these subproblems every time dynamics occur [53]. GND efficiently allocates requests to neighborhoods of agents based on the orbital geometry of the constellation and has been shown to effectively partition COSP instances. For completeness,

Algorithm 1 Dynamic Incremental Neighborhood Stochastic Search (D-NSS) for agent a

Input: $sched, R, A, S_a, C_a, \Upsilon, maxIters$

Output: Schedule for agent a

- 1: $\mathcal{N} = \text{COMPUTESUBPROBLEM}(a, A, R, S_a, \Upsilon)$
 - 2: $sched = \text{REPAIR}(sched, R_{\mathcal{N}}, S_a, C_a)$
 - 3: $sched = \text{D-NSS-SEARCH}(sched, R_{\mathcal{N}}, A_{\mathcal{N}}, S_a, C_a, maxIters)$
 - 4: **return** $sched$
-

Algorithm 2 REPAIR for agent a

Input: $sched, R, S_a, C_a$ **Output:** Repaired Schedule for agent a

```
1: for  $s \in sched$  do
2:   if  $r(s) \notin R$  then
3:     REMOVEFROMSCHEDULE( $sched, s$ )
4:   shuffle  $S_a$ 
5:   for  $s \in S_a$  do
6:     if  $s$  satisfies  $C_a \wedge r(s) \notin sched$  then
7:       ADDTOSCHEDULE( $sched, s$ )
8:   return  $sched$ 
```

Algorithm 3 D-NSS-SEARCH for agent a

Input: $sched, R_N, A_N, S_a, C_a, maxIters$ **Output:** Schedule for agent a

```
1: while  $i < maxIters \wedge$  not converged do
2:    $com\_out, com\_in = MESSAGE(A_N, sched)$ 
3:   shuffle  $R_N$ 
4:   for  $r \in R_N$  do
5:     assigned = STOCHASTICUPDATE( $r, sched, com\_in$ )
6:     if assigned = TRUE then
7:        $sched = SCHEDULE(r, sched, S_a)$ 
8:    $i \leftarrow i + 1$ 
9:   return  $sched$ 
```

we provide an outline of GND. GND was first introduced with the NSS algorithm and we refer the reader to prior work for a full presentation of the heuristic [53].

GND leverages the orbital geometry of the satellite constellation to hierarchically partition requests to neighborhoods of agents. Let K be the set of orbital planes that define a satellite constellation¹. For every request, an agent computes the supply from all orbital planes and adds these to estimate the total number of agents with overflights for the request. The supply can also be thought of as an estimate of degrees in the constraint graph. Iterating through requests in ascending order of supply, a request gets assigned to the n neighborhoods with the highest ratio of supply to temporal conflicts. Temporal conflicts are counts of other requests already allocated to a neighborhood that overlap in time with a given request. Finally, within a neighborhood, requests are further subdivided to agents based on biases towards specific tiles on Earth. GND with n degrees of incompleteness is denoted GND(n).

D-NSS restarts the search phase when changes are initiated in the problem. A key procedure is the REPAIR function that repairs previously computed solutions to leverage assignments of unchanging variables shown in Algorithm 2. Repairing consists of removing all tasks that are no longer in the current problem instance and greedily inserting new tasks in random order. Note that we can also remove from an agent’s schedule all requests that have been previously executed by agents in the same neighborhood. One benefit of the D-NSS algorithm is that it can leverage different repair procedures. We use the random repair function for two main reasons.

¹An orbital plane is a fixed orbit around Earth that many satellites may follow at various spacing.

Random initialization is the standard for DSA and NSS variants in static domains [49] as it promotes diverse solutions and D-NSS is designed to be as lightweight as possible. Computation-intensive repair procedures counteract the efficiency of the algorithm. We show later on that performing random repair improves the quality and efficiency of D-NSS on DCOSP instances.

After each agent repairs its solution, all agents synchronously begin the iteration phase, D-NSS-SEARCH, to fine-tune the repaired solutions. D-NSS-SEARCH extends the search phase of NSS to account for tasks that have just been scheduled versus ones that have been executed. We detail the following sub-procedures.

- COMPUTESUBPROBLEM(a, A, R, S_a, Y). This function computes the neighborhood of agent a and the subset of requests for that neighborhood using the decomposition heuristic Y . We use $Y = \text{GND}$. GND produces neighborhoods that are reflexive and transitive.
- MESSAGE($A_N, sched$). This function defines the message exchange between agents in a neighborhood. Each agent a sends to each other agent in $A_N \setminus \{a\}$ the subset of R_N that it has executed a task for already and the subset that it has scheduled in the previous iteration via the variable com_out . The resulting data structure com_in contains the satisfaction information for the neighborhood.
- STOCHASTICUPDATE($r, sched, com_in$). This function computes the assignment of an agent and a request based on the neighborhood’s communication. An assignment refers to if an agent should attempt to schedule a specific request. Let W be the count of agents that scheduled or executed r in the previous iteration, the probability P_u be a hyperparameter, executed(r) be the predicate that r was executed already, and assigned(a, r) be the predicate that a is assigned to r . An agent computes the probability of assigning to r in the next iteration, $P(a, r | com_in)$ using the update scheme from Table 1. For example, according to com_in , if request r has not been executed, agent a is not assigned to it, and $W \geq 1$, agent a will always remain unassigned to r . This update scheme extends the static NSS update schemes to account for dynamic scheduling and execution. If a task for a request has already been executed, then all agents should unassign. Note that executed(r) $\implies W \geq 1$.
- SCHEDULE($r, sched, S_a$). This function tries to schedule a task for request r given the current schedule. If a task for r satisfies C_a it is inserted into the schedule. Otherwise, the scheduler may remove a single task from the schedule to satisfy the constraints. The task with the closest start time to the new task is used as a heuristic for removal. Agent a remains

	executed(r)	\neg executed(r) \neg assigned(a, r)	\neg executed(r) assigned(a, r)
$W = 0$	N/A	1	$1 - P_u$
$W \geq 1$	0	0	$1/W$

Table 1: Stochastic assignment update scheme for agent a . The table values denote the probability that agent a assigns to request r based on $com_in, P(a, r | com_in)$.

assigned to a removed task and can attempt to re-schedule it in subsequent iterations. Task removal enables the search phase to overcome local maxima. Note that vanilla DCOSP considers all requests of equal priority and there is no temporal flexibility in the start or end times of tasks. However, DCOSP and the scheduling procedure are amenable to these extensions.

Without prior knowledge of the problem dynamics, it is difficult for any online algorithm to reason about which scheduled tasks will be executed. Prioritizing requests that are earlier in the horizon may be more likely not to be removed and yield reward. However, expending resources early on in the horizon results in fewer resources available to handle problem dynamics later. Although D-NSS currently focuses on reactive repair, it can be extended to integrate proactive scheduling through predictive models, which is a subject of future work.

We introduce variables to analyze the complexity of D-NSS. Let L be the maximum size of a satellite’s schedule, $A_{\mathcal{N}}$ be the largest set of agents in a sub-problem, and $R_{\mathcal{N}}$ be the largest set of requests in a sub-problem. In general $L \ll |R_{\mathcal{N}}| \ll |R|$ since L is constrained by resources and time. Sub-problem computation via GND has a time complexity of $O(|R_{\mathcal{N}}|)$ and uses no communication. The REPAIR procedure is individually computed by an agent and takes $O(|R_{\mathcal{N}}| \log L)$ time and again uses no communication. D-NSS inherits from NSS a computation and communication complexity of $O(|A_{\mathcal{N}}| \cdot |R_{\mathcal{N}}|)$ during an iteration.

4.3 Baseline Algorithms

We evaluate several baseline algorithms in addition to D-NSS. The naive alternative to D-NSS is to recompute a solution from scratch every time problem dynamics occur. We refer to this procedure as 0-NSS. 0-NSS has the same theoretical complexities as D-NSS. We also evaluate dynamic variations of the *Distributed Stochastic Search Algorithm* (DSA) [33, 49]. D-DSA uses Algorithm 2 to repair DSA solutions. Likewise, 0-DSA runs DSA from scratch every time the problem changes. These DSA variants have a complexity of $O(|A| \cdot |R|)$ per iteration. Dynamic DSA is one of the most lightweight DDCCOP solvers and can be deployed without violating the constraint graph assumptions of DCOSP. Other DDCCOP algorithms such as a dynamic variation of MGM [26], would also incur a computation and communication complexity of $O(|A| \cdot |R|)$ per iteration. Due to the scale of the problem instances and the constraints of satellite computation and communication, we are unable to evaluate any exponential-time algorithms.

We include two baseline non-communication-reliant algorithms: random and greedy. Both of these algorithms construct a schedule for a satellite without reasoning about other agents. These algorithms transition schedules after the dynamics occur by removing redundant tasks. The random solver randomly inserts tasks into an agent’s schedule while the greedy algorithm orders tasks by their start time and iteratively inserts them into a schedule in a single pass. Although simple, a greedy solver is comparable to deployed planners on spacecraft [15].

Finally, we use the construction outlined previously to obtain optimal and near-optimal solutions. For small problems, we obtain an optimal solution by constructing a static COSP instance and solving

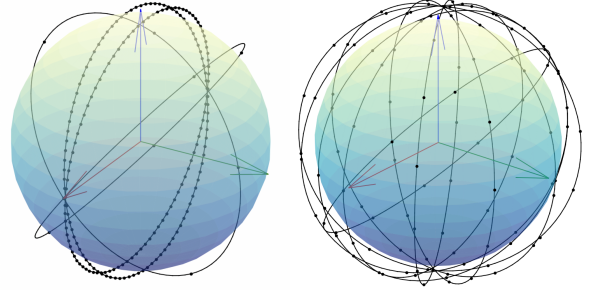


Figure 2: The satellite constellations: Planet (left) and Walker (right). Dots show satellites in an orbital plane.

it with a centralized branch and bound. For large problems, computing an optimal solution is not computationally feasible, so we lower bound the optimal solution using *Squeaky Wheel Optimization* [21], an incomplete centralized solver.

5 EXPERIMENTS

5.1 Setup

In this section, we outline the experimental setup including constellations and dynamic observation campaigns.

5.1.1 Satellite Constellations. We evaluate two constellations modeled after operational low Earth orbit constellations [38]. The Planet constellation is modeled on the Dove constellation from Planet Labs. This constellation is composed of two near sun-synchronous orbital planes at 95° inclinations each composed of 95 satellites with an additional two orbital planes at 52° inclinations each with 5 satellites. The Walker constellation is motivated by the Skysat constellation from Planet Labs. This constellation has 6 orbital planes with 14 satellites each at an 88° inclination and an overlay of 2 orbital planes at a 51.6° inclination with 12 satellites. Each satellite has a memory capacity of 125 GB and a single sensor that can slew to 60° and 45° off-nadir for the Planet and Walker constellations respectively. Figure 2 depicts these two constellations.

5.1.2 Downlinks. A 62.5 MB/s constant bit stream models the satellite downlink during visibility periods with a ground station. We incorporate two ground stations: the ASF Near Space Network Satellite Tracking Ground Station and the Guam Remote Ground Terminal System.

5.1.3 Dynamic Observation Request Campaigns. The target set, \mathcal{T} , is 634 globally distributed cities. Dynamic campaigns consist of periodic requests of these targets. For large problem instances, a periodicity is uniformly sampled from the range [5, 12]. A target with periodicity p is requested to be observed once within p evenly spaced intervals. For small problem instances, we fix the periodicity at 3. The scheduling horizon is set to be one day. However, the start of the horizon is randomly initialized. An observation’s memory consumption is sampled from a normal distribution with mean 50 MB and standard deviation 10 MB. The time interval required to execute an observation is 63 seconds which accounts for imaging, slewing, and processing.

To generate dynamics, we select a volatility parameter v . This determines the number of changes during the horizon. The time of the dynamics is uniformly distributed over the final $1 - \frac{2}{3 \cdot v}$ of the scheduling horizon. This ensures that dynamics do not occur directly after the start of the horizon. The set of starting requests is randomly selected from the request set and initialized to one-third of the total size. We select from the remaining requests a random set of $\frac{2}{3 \cdot v}$ to add. From the active requests, we randomly remove $\frac{1}{3 \cdot v}$ of them. We enforce that once a request is removed, it is not added back, and requests are not changed after their execution horizon starts. We sample $v \in [3, 5]$ uniformly.

5.1.4 Hyperparameters and Execution Environment. The stochastic update of NSS and DSA relies on the probability, P_u that determines when an agent should unassign from a task. We fix this value at 0.7, as published in previous work [33, 53]. The *maxIters* is set to 20. For the GND heuristic, we set all hyperparameters as described in previous work, and specifically use the GND(2) heuristic [53]. We use a random seed of 2005 for the initial scenario generation. When evaluating N scenarios, we increment this seed for generation of each subsequent simulation. The random seed used in the REPAIR procedure is initialized to 1. The NSS and DSA algorithms are seeded with a value of 1234, and the random seed for GND is set to 2. The random scheduling algorithm uses a seed of 2023. All experiments are executed using Java 19 on a MacBook Pro 16 laptop with an M2 Max processor (12-core CPU and 38-core GPU) and 64 GB of RAM.

5.2 Results on Small Problem Instances

We are able to obtain an optimal solution for small problem instances using the omniscient offline algorithm. We solve this DCOP with a centralized branch and bound to obtain an optimal schedule for each satellite. For the Planet constellation, we solve campaigns

Algorithm	Opt. Gap (%)	Time (ms)	Messages (KB)
Random	2.530	<1	0
Greedy	8.373	<1	0
D-NSS	1.867	<1	7.3
0-NSS	2.590	<1	13.2
D-DSA	4.217	5.3	980.6
0-DSA	3.795	4.7	718.4

Table 2: Results on 10 dynamic problems for the Planet constellation (up to 500 requests).

Algorithm	Opt. Gap (%)	Time (ms)	Messages (KB)
Random	14.945	<1	0
Greedy	15.604	<1	0
D-NSS	0.142	5.2	240.2
0-NSS	0.480	6.1	400.0
D-DSA	1.165	58.2	10,459.5
0-DSA	1.215	54.0	7,789.0

Table 3: Results on 10 dynamic problems for the Walker constellation (up to 1000 requests).

of up to 500 requests. For the Walker constellation, this increases to 1000 requests. In addition to having fewer agents, the Walker constellation geometry under-constrains small problems, making finding optimal solutions faster. An optimal solver does not consistently terminate for larger problems. We report the average gap in satisfaction percentage to the optimal solution for 10 dynamic small problem instances in Tables 2 and 3.

These results support the theoretical analysis of D-NSS and demonstrate near-optimal performance. D-NSS outperforms all baselines, including D-DSA, 0-DSA, and 0-NSS, and does so while using less computation and communication. Notably, compared to DSA variants, D-NSS finds better solutions using an order of magnitude less computation and up to two orders of magnitude less communication. This is due to both the lower theoretical complexity per iteration and the faster convergence of D-NSS.

5.3 Results on Large Problem Instances

We evaluate the algorithms on realistic, large-scale, dynamic scenarios with thousands of requests. Figure 3 shows the results for the Planet and Walker constellations. We report the total utility of each algorithm in Figures 3a and 3d, the total message volume in Figures 3b and 3e, and the average per-agent execution time in Figures 3c and 3f. Note the log scale in the figures for message volume and runtime.

In terms of solution quality, D-NSS and D-DSA achieve close to the optimal lower-bound solution. D-NSS outperforms both 0-NSS and 0-DSA as well as the greedy and random baselines. Crucially, D-NSS achieves high solution quality while having a lower total message volume and execution time compared to D-DSA, 0-DSA, and 0-NSS. Both D-NSS and 0-NSS use an order of magnitude fewer messages and computation times than their DSA counterparts. Since the optimal lower bound algorithm is computed centrally, the message volume corresponds to the schedules a ground station would have to uplink.

We also evaluate the stability and convergence of D-NSS compared to D-DSA, 0-DSA, and 0-NSS. Figures 3g and 3h show the average solution quality over iterations of the algorithms during large-problem instances with $v = 5$. We fix the number of iterations for all algorithms to show the relative performance. Clearly, D-NSS and D-DSA are much more stable, reacting much less volatility to problem dynamics, illustrated by the sharp drops in solution quality of 0-NSS and 0-DSA. Notably, D-NSS is very stable; even in the iteration directly after problem dynamics, D-NSS repairs solutions effectively to maintain or improve solution quality. This supports computing and repairing sub-problems rather than global solutions like D-DSA. In addition, the stability of D-NSS leads to a quick convergence in practice, which drives the efficiency of the algorithm.

6 CONCLUSION

Large problems remain a challenge for DDCOP algorithms due to their computational and communication complexities. Even linear time and messaging solutions may not be feasible when constraint graphs are highly connected or agents have limited compute. These challenges apply to DCOSP. We present the D-NSS algorithm, a decomposition-based algorithm that is efficient in time and message

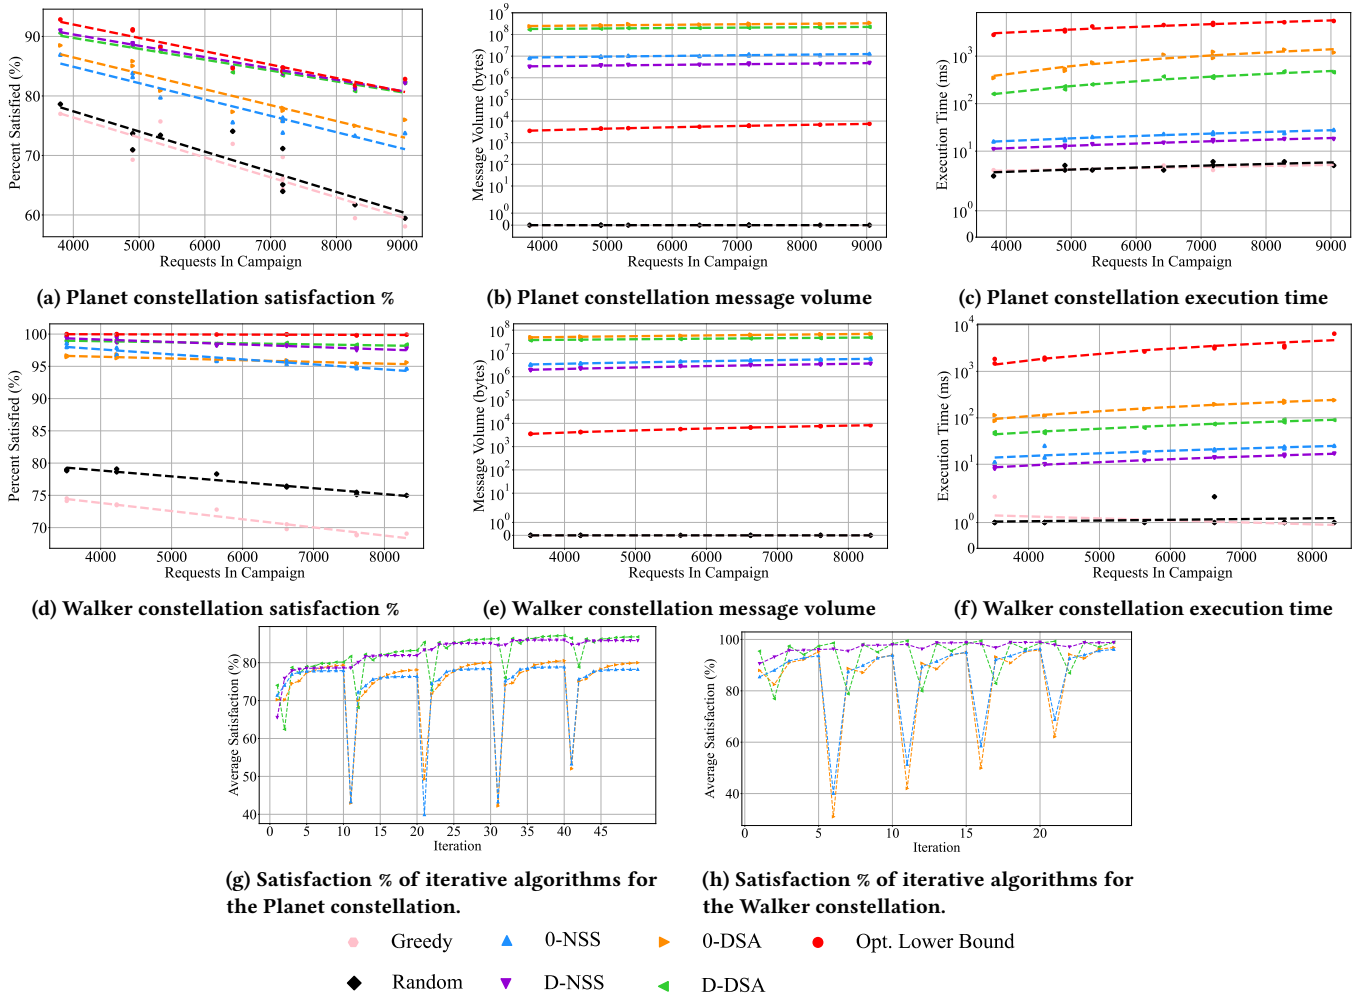


Figure 3: Results of 10 large-scale simulations for the Planet (top row) and Walker (middle row) constellation. We report the average satisfaction, total message volume, and per-agent runtime. Note the log scale for message volume and runtime. We also show the solution quality across a fixed number of iterations for iterative algorithms for instances with $v = 5$ (bottom).

complexity and can scale to problems much larger than previous approaches. We show that D-NSS stabilizes quickly and uses an order of magnitude fewer messages and compute times compared to DDCOP baselines. Despite no quality guarantees, we show that D-NSS computes near-optimal solutions.

D-NSS generalizes to many DDCOP problems. While we leverage the GND heuristic, the D-NSS framework can accommodate alternative heuristics, including learned or adaptive ones, enabling broader applicability across DDCOP domains. Some examples include multi-agent path finding [42], mobile sensor teams [34], and UAV coordination [39].

In future work, we can extend D-NSS to reason about future dynamics. We can construct or learn heuristics that prioritize tasks that are more likely to yield utility when executed. Alternatively, we can extend DCOSP to be a *Proactive DCOP* (PDDCOP) [20] where we have priors on problem dynamics, and agents can generate

robust offline solutions. However, PDDCOPs require a model of problem dynamics which may not be obtainable in practice.

DCOSP solutions will be demonstrated in operational scenarios beginning in 2026. One example is NASA’s FAME demonstration [9] which centers around a *federated* observation system [50]. The FAME demonstration consists of more than 60 participating Earth-observing spacecraft that will coordinate their measurements to optimize observation completion across dynamic scenarios. Through onboard control, satellites will be able to react faster and with more precision to time-sensitive events such as natural disasters. These observations will provide novel measurements of scientific processes and crucial, up-to-date information for human operators. This work has laid the foundation for dynamic, large-scale, distributed onboard scheduling of Earth-observing satellites.

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 under Grant No. 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] Sean Augenstein, Alejandra Estanislao, Emmanuel Guere, and Sean Blaes. 2016. Optimal scheduling of a constellation of Earth-imaging satellites, for maximal data throughput and efficient human management. In *Proceedings of the International Conference on Automated Planning and Scheduling*, Vol. 26. 345–352.
- [2] Graham Billiau, Chee Fon Chang, and Aditya Ghose. 2010. SBDO: A new robust approach to dynamic distributed constraint optimisation. In *International Conference on Principles and Practice of Multi-Agent Systems*. 11–26.
- [3] James Boerkoel, James Mason, Daniel Wang, Steve Chien, and Adrien Maillard. 2021. An efficient approach for scheduling imaging tasks across a fleet of satellites. In *Proceedings of the International Workshop on Planning and Scheduling for Space*.
- [4] Grégory Bonnet and Catherine Tessier. 2007. Collaboration among a satellite swarm. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems*. 1–8.
- [5] Grégory Bonnet and Catherine Tessier. 2008. Coordination despite constrained communications: A satellite constellation case. In *Proceedings of the National Conference on Control Architectures of Robots*. 89–100.
- [6] Abhijit Chatterjee and Ratnasingham Tharmarasa. 2024. Multi-stage optimization framework of satellite scheduling for large areas of interest. *Advances in Space Research* 73, 3 (2024).
- [7] Steve Chien, Alberto Candela, Itai Zilberstein, David Rijlaarsdam, Tom Hendrix, and Aubrey Dunne. 2024. Leveraging commercial 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*.
- [8] 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.
- [9] Steve Chien, Itai Zilberstein, Alberto Candela, Domenico Barretta, David Rijlaarsdam, Tom Hendrix, Aubrey Dunne, Oriol Cortés Grauc, Alexandre Gol i Mestre, 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*.
- [10] Steve Chien, Itai Zilberstein, Alberto Candela, David Rijlaarsdam, Amaury Perrocheau, Aubrey Dunne, Tom Hendrix, Oriol Cortés Grauc, Alexandre Gol i Mestre, 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*.
- [11] Duncan Eddy and Mykel J Kochenderfer. 2021. A maximum independent set method for scheduling Earth-observing satellite constellations. *Journal of Spacecraft and Rockets* 58, 5 (2021), 1416–1429.
- [12] Adrian Petcu and Boi Faltings. 2005. Superstabilizing, fault-containing multiagent combinatorial optimization. In *AAAI* 449–454.
- [13] Ferdinando Fioretto, Enrico Pontelli, and William Yeoh. 2018. Distributed constraint optimization problems and applications: A survey. *Journal of Artificial Intelligence Research* 61 (2018), 623–698.
- [14] Ferdinando Fioretto, William Yeoh, Enrico Pontelli, Ye Ma, and Satishkumar J Ranade. 2017. A distributed constraint optimization (DCOP) approach to the economic dispatch with demand response. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems*. 999–1007.
- [15] Dan Gaines, Steve Chien, Gregg Rabideau, Stephen Kuhn, Vincent Wong, Amruta Yelamanchili, Shannon Towey, Jagriti Agrawal, Wayne Chi, Andrea Connell, Evan Davis, and Colette Lohr. 2022. Onboard planning for the Mars 2020 Perseverance Rover. In *Symposium on Advanced Space Technologies in Robotics and Automation*.
- [16] Amir Gershman, Amnon Meisels, and Roie Zivan. 2009. Asynchronous forward bounding for distributed COPs. *Journal of Artificial Intelligence Research* 34 (2009), 61–88.
- [17] Al Globus, James Crawford, Jason Lohn, and Anna Pryor. 2004. A comparison of techniques for scheduling Earth-observing satellites. In *Proceedings of the Conference on Innovative Applications of Artificial Intelligence*.
- [18] Lei He, Xiaolu Liu, Gilbert Laporte, Yingwu Chen, and Yingguo Chen. 2018. An improved adaptive large neighborhood search algorithm for multiple agile satellites scheduling. *Computers & Operations Research* 100 (2018), 12–25.
- [19] Katsutoshi Hirayama and Makoto Yokoo. 1997. Distributed partial constraint satisfaction problem. In *Proceedings of the International Conference on Principles and Practice of Constraint Programming*. 222–236.
- [20] Khoi D Hoang, Ferdinando Fioretto, Ping Hou, William Yeoh, Makoto Yokoo, and Roie Zivan. 2022. Proactive dynamic distributed constraint optimization problems. *Journal of Artificial Intelligence Research* 74 (2022), 179–225.
- [21] David E Joslin and David P Clements. 1999. Squeaky wheel optimization. *Journal of Artificial Intelligence Research* 10 (1999), 353–373.
- [22] Sankalp Khanna, Abdul Sattar, David Hansen, and Bela Stantic. 2009. An efficient algorithm for solving dynamic complex DCOP problems. In *International Joint Conference on Web Intelligence and Intelligent Agent Technology*, Vol. 2. 339–346.
- [23] Parimal Kopardekar and Laurie Grindle. 2021. NASA ARMD Wildfire Management Workshop. (2021).
- [24] Robert N Lass, Evan Sultanik, and William C Regli. 2008. Dynamic distributed constraint reasoning. In *AAAI* 1466–1469.
- [25] Xiang Lin, Yuning Chen, Junhua Xue, Boquan Zhang, Lei He, and Yingwu Chen. 2024. Large-volume LEO satellite imaging data networked transmission scheduling problem: Model and algorithm. *Expert Systems with Applications* 249 (2024), 123649.
- [26] Rajiv T Maheswaran, Jonathan P Pearce, and Milind Tambe. 2004. Distributed algorithms for DCOP: A graphical-game-based approach. In *Proceedings of the International Conference on Parallel and Distributed Computing Systems*. 432–439.
- [27] Roger Mailler. 2005. Comparing two approaches to dynamic, distributed constraint satisfaction. In *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems*. 1049–1056.
- [28] Roger Mailler and Victor Lesser. 2004. Solving distributed constraint optimization problems using cooperative mediation. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems*. 438–445.
- [29] Pragnesh Jay Modi, Wei-Min Shen, Milind Tambe, and Makoto Yokoo. 2005. ADOPT: Asynchronous distributed constraint optimization with quality guarantees. *Artificial Intelligence* 161, 1-2 (2005), 149–180.
- [30] Sreeja Nag, Alan S Li, and James H Merrick. 2018. Scheduling algorithms for rapid imaging using agile Cubesat constellations. *Advances in Space Research* 61, 3 (2018), 891–913.
- [31] NewSpace. 2023. NewSpace Constellations. <https://www.newspace.im>. Accessed: 2025-05-01.
- [32] Duc Thien Nguyen, William Yeoh, Hoong Chuin Lau, and Roie Zivan. 2019. Distributed Gibbs: A linear-space sampling-based DCOP algorithm. *Journal of Artificial Intelligence Research* 64 (2019), 705–748.
- [33] Shreya Parjan and Steve A. Chien. 2023. Decentralized observation allocation for a large-scale constellation. *Journal of Aerospace Information Systems* (2023), 1–15.
- [34] Arseniy Pertzovsky, Roie Zivan, and Noa Agmon. 2024. Collision avoiding MaxSum for mobile sensor teams. *Journal of Artificial Intelligence Research* 79 (2024), 1281–1311.
- [35] Adrian Petcu and Boi Faltings. 2005. DPOP: A scalable method for multiagent constraint optimization. In *Proceedings of the International Joint Conference on Artificial Intelligence*. 266–271.
- [36] Sean Phillips and Fernando Parra. 2021. A case study on auction-based task allocation algorithms in multi-satellite systems. In *Proceedings of AIAA Scitech*.
- [37] Gauthier Picard. 2022. Auction-based and distributed optimization approaches for scheduling observations in satellite constellations with exclusive orbit portions. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems*. 1056–2064.
- [38] Planet. 2023. Our Constellations. <https://www.planet.com/our-constellations>. Accessed: 2025-05-01.
- [39] Marc Pujol-Gonzalez, Jesus Cerquides, Pedro Meseguer, Juan Antonio Rodríguez-Aguilar, and Milind Tambe. 2013. Engineering the decentralized coordination of UAVs with limited communication range. *Advances in Artificial Intelligence* 1 (2013), 199–208.
- [40] 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*.
- [41] Anton Ridgway and Roger Mailler. 2015. Dynamic theoretical analysis of the distributed stochastic and distributed breakout algorithms. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems*. 405–412.
- [42] 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.

- [43] Vishwa Shah, Vivek Vittaldev, Leon Stepan, and Cyrus Foster. 2019. Scheduling the world's largest Earth-observing fleet of medium-resolution imaging satellites. In *Proceedings of the International Workshop on Planning and Scheduling for Space*. 156–161.
- [44] Samuel Squillaci, Cédric Pralet, and Stéphanie Roussel. 2023. Scheduling complex observation requests for a constellation of satellites: Large neighborhood search approaches. In *Proceedings of the International Conference on the Integration of Constraint Programming, Artificial Intelligence, and Operations Research*. 443–459.
- [45] Samuel Squillaci, Stéphanie Roussel, and Cédric Pralet. 2021. Managing complex requests for a constellation of Earth-observing satellites. In *Proceedings of the International Workshop on Planning and Scheduling for Space*.
- [46] Ruben Stranders, Alessandro Farinelli, Alex Rogers, and Nick Jennings. 2009. Decentralised coordination of mobile sensors using the max-sum algorithm. In *Proceedings of the International Joint Conference on Artificial Intelligence*. 299–304.
- [47] Xinwei Wang, Guohua Wu, Lining Xing, and Witold Pedrycz. 2020. Agile Earth observation satellite scheduling over 20 years: Formulations, methods, and future directions. *IEEE Systems Journal* 15, 3 (2020), 3881–3892.
- [48] William Yeoh, Pradeep Varakantham, Xiaoxun Sun, and Sven Koenig. 2015. Incremental DCOP search algorithms for solving dynamic DCOP problems. In *International Conference on Web Intelligence and Intelligent Agent Technology*, Vol. 2. 257–264.
- [49] 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.
- [50] 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*.
- [51] 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*.
- [52] Itai Zilberstein and Steve Chien. 2026. Large-Scale Continual Scheduling and Execution for Dynamic Distributed Satellite Constellation Observation Allocation. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems*. Extended Abstract.
- [53] 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.
- [54] Roie Zivan, Harel Yedidsion, Steven Okamoto, Robin Grinton, and Katia Sycara. 2015. Distributed constraint optimization for teams of mobile sensing agents. *Autonomous Agents and Multiagent Systems* 29 (2015), 495–536.