

Going Beyond Mono-Mission Earth Observation: Using the Multi-Agent Paradigm to Federate Multiple Missions

Blue Sky Ideas Track

Jean-Loup Farges
ONERA/DTIS, Université de Toulouse
Toulouse, France
jean-loup.farges@onera.fr

Filipo Perotto
ONERA/DTIS, Université de Toulouse
Toulouse, France
filipo.perotto@onera.fr

Gauthier Picard
ONERA/DTIS, Université de Toulouse
Toulouse, France
gauthier.picard@onera.fr

Cédric Pralet
ONERA/DTIS, Université de Toulouse
Toulouse, France
cedric.pralet@onera.fr

Cyrille de Lussy
Airbus Defence and Space
Toulouse, France
cyrille.de-lussy@airbus.com

Jonathan Guerra
Airbus Defence and Space
Toulouse, France
jonathan.guerra@airbus.com

Philippe Pavero
Airbus Defence and Space
Toulouse, France
philippe.pavero@airbus.com

Fabrice Planchou
Airbus Defence and Space
Toulouse, France
fabrice.planchou@airbus.com

ABSTRACT

We envision a multi-agent federation layer to coordinate systems composed of independent Earth observation missions. The goal of this federation is to allow clients requesting acquisitions of large areas to easily access several constellations of satellites and communication sites to compose and download their acquisitions, in a reduced time compared to conventional uncoordinated requests. We identify several scientific tracks and challenges related to agent-based approaches such as coordination, planning and learning, to implement to two key federation functions: (i) multi-mission coverage feasibility and dispatching and (ii) communication site booking.

KEYWORDS

Earth observation, mission federation, coverage, communication

ACM Reference Format:

Jean-Loup Farges, Filipo Perotto, Gauthier Picard, Cédric Pralet, Cyrille de Lussy, Jonathan Guerra, Philippe Pavero, and Fabrice Planchou. 2024. Going Beyond Mono-Mission Earth Observation: Using the Multi-Agent Paradigm to Federate Multiple Missions: Blue Sky Ideas Track. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024)*, Auckland, New Zealand, May 6 – 10, 2024, IFAAMAS, 5 pages.

1 INTRODUCTION

Earth Observation (EO) market demands faster and more efficient satellite systems to meet the growing needs for responsive, revisiting, and multi-mission¹ capabilities [10]. Applications such as

¹A mission typically consists in a single satellite or a constellation of satellites managed by a single operator and a set of communication assets to uplink and downlink data.



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), N. Alechina, V. Dignum, M. Dastani, J.S. Sichman (eds.), May 6 – 10, 2024, Auckland, New Zealand. © 2024 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

border surveillance, maritime monitoring, and disaster response require rapid satellite programming to provide timely imagery. Moreover, users seek scalable and cost-effective solutions to support their environmental and agricultural sustainability efforts. To address these demands, new EO systems are equipped with advanced instruments and satellites, enabling high resolution, large memory, great acquisition capacity, and fast image acquisition across multiple locations, resulting in significant performance improvement.

However, users and satellite operators are demanding more responsive and agile systems. Traditional EO systems, often consisting of just one or two satellites, fall short with sufficient revisit capabilities and reactivity in image delivery. Users require access to multiple acquisition means, including other satellites, instruments, and even non-EO missions, to effectively address their needs. In this paper, we particularly focus on the upstream of the ground segment, which encompasses satellite control, mission programming, and image validation, that are mission-specific and tightly integrated with the satellite system. This siloed approach (see Figure 1) is optimized for individual missions and becomes problematic when managing multiple EO systems. The lack of interoperability between these systems impedes efficient utilization of their overall capacity.

To answer the strong need for a unified solution that enables seamless access to various information systems, the DOMINO-E project [9] aims to develop a federation layer (see Figure 2) that empowers EO operators to coordinate and prioritize their imagery needs across different systems, using the multi-agent paradigm. While multi-agent systems have already been identified as a relevant approach for EO constellations [18], this paper proposes a novel vision to orchestrate multiple constellations. The proposed architecture's modular, flexible, and resilient design should allow for the integration of additional relevant services without compromising system performance. We particularly address challenges related to two key services for multi-mission operation:

- *Multi-Mission Coverage Feasibility and Dispatch* (Section 2): in case of extended surface coverage to be performed in a

given period, one mission is often not enough to accomplish the task on time and on quality. Thus, we propose a multi-agent system that enables dynamic selection among several satellite systems for coverage.

- *Multi-Mission Communication Booking* (Section 3): ground systems currently rely on dedicated antennas for satellite control, data transfer, and image downloads. The choice of antennas is typically fixed throughout the system’s lifespan. We propose a dynamic antenna reservation system, leveraging Ground Segment as a Service (GSaaS), to ensure adequate data downlink capacity and optimize costs.

Centralized solutions might not work to implement these complex services because key information is spread out. So, we need to install cooperation between components and stakeholders, suggesting techniques from multi-agent planning, resource allocation, auctions, and learning will be crucial for its development.

2 FEDERATED OBSERVATION

This section addresses the following decision problem: "How to divide a large area and to assign these subdivisions to different missions in order to minimize time of acquisition and maximize the quality of the images, even though we do not know the future workload of missions?" Such decision has to be made each time a new request is received or when a mission updates its dates of acquisition.

2.1 System Composition

The problem considered involved a set of K end-users, referred to as *clients*, who request images over large areas on the Earth’s surface. It also involves a set of N independent observation agents, referred to as *missions*, developed over the years to fulfill specific observation needs. Basically, each mission disposes of a set of low-Earth orbit satellites capable of collecting images on elementary meshes, where each mesh usually covers a few tens of square kilometers, and each mission disposes of its own *planner* to determine the activities of its satellites depending on the observation requests it receives.

The global objective is then to complete the observations requested by the K clients using the N observation missions available. For this, the conventional approach is that each client posts its observation requests directly to the mission of its choice (see Fig. 1). To go beyond this paradigm, the *federation layer* will receive the observation requests of the clients, dispatch observation tasks to the missions, get the observation data back, and deliver the images to the clients (see Fig. 2). Thus, each client benefits from a seamless access to numerous satellite resources without having to care about the details. Also, the federation layer can exploit for the best the resources available to significantly speed up the observation of a large area by dividing it into several sub-areas that can be observed in parallel by distinct missions, and merge observation requests over overlapping areas to reduce the global load.

2.2 Multi-Agent Decisions and Problems

To define the federation layer, the first point consists in defining how the different agents interact. A first approach would be to design an architecture implementing *multi-agent negotiation* [12], where the federation layer would send queries to missions while

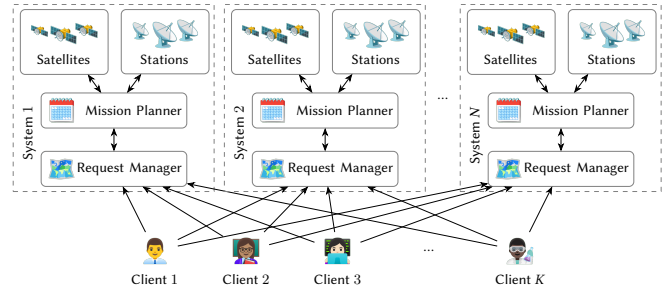


Figure 1: The conventional architecture is vertical and users must interrogate and send requests to each system separately.

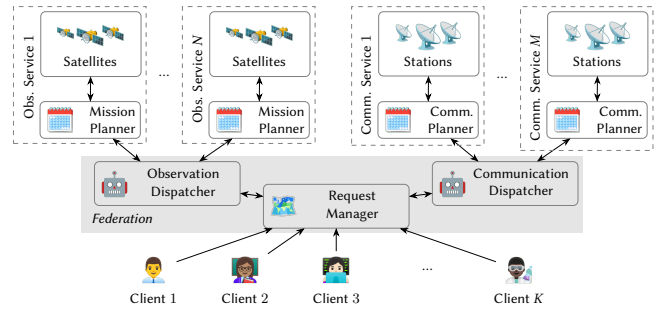


Figure 2: The DOMINO-E architecture calls different systems as services, optimizing combined requests from multiple clients to multiple systems in a transparent way.

the latter would accept or not to make observations over some sub-areas. Then, each sub-area would be actually allocated to a specific mission. But, such a negotiation scheme is not fully compatible with the presence of clients that are waiting for images, especially if the response time of a mission is high during the negotiation part. Indeed, assessing the impact of including a new observation into a mission schedule is computer intensive, or may require an efficient surrogate model [22]. To limit the number of interactions between the agents, an alternative could be implementing *combinatorial auctions* [7], involving a first phase where the federation layer asks the missions to bid on possible observation sub-areas, and a second phase where observation tasks are dispatched given the bids received. But such an approach would require developing a bidding system for each mission, which is time-consuming.

This is why we envision an architecture where the federation layer skips the bidding part and directly allocates observation tasks based on its current knowledge of the capacity and load of each mission. In this case, the communication with the individual missions is minimal, but the definition of the envisioned federation layer raises several challenges detailed thereafter.

2.3 Challenges for the Federation Layer

Highly combinatorial problems. To make dispatching decisions, the federation layer must tackle a highly combinatorial problem (hundreds or thousands of requests, tens or hundreds of satellites, numerous ways to partition the area of each request into a set of sub-areas, etc.). On this point, there is a need to study both coarse-grain dispatching strategies where large sub-areas are allocated to

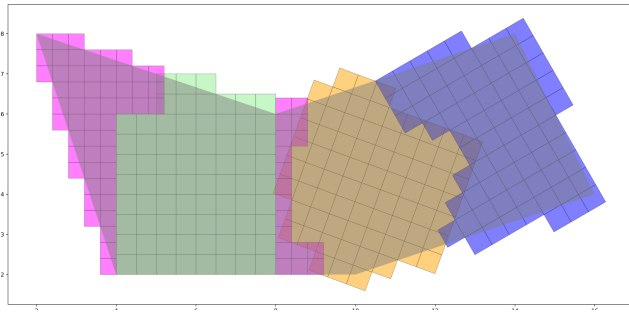


Figure 3: Multi-mission coverage of a large area (in grey) using meshes of four different missions (one color per mission)

the missions, and fine-grain dispatching strategies where detailed meshes are considered. Reasoning at the mesh level may lead to better dispatching decisions but increases the difficulty of the optimization problem to solve. An example of a possible multi-mission coverage strategy defined at the mesh level is illustrated in Fig. 3. The federation layer must also handle multiple objectives, including the minimization of the completion time of the coverage, and the minimization of the total load of the system. In another direction, the federation layer must take into account the compatibility between the requirements associated with each request and the capabilities of the agents available, since the satellites of two distinct missions can be equipped with instruments that have different resolutions or spectral features.

Need for a model of the other agents. To allocate observation tasks, the federation layer must handle a model of the current capabilities of each mission. One difficulty is that such a model can be imprecise, especially for *external missions* whose resources are not under the direct control of the system manager. And even for the *legacy missions*, the federation layer does not have a full control due to the presence of a specific planning engine associated with each mission. This is why the federation layer should ideally *learn a high-level model of the capabilities* of the missions. Also, each individual mission can receive urgent high-priority requests that may postpone the achievement of low-priority observations over large areas, and for this the federation layer should consider exploiting a model of the density of the high-priority requests within a given area, based for instance on historical request data.

Uncertainty management. In the problem considered, there are various sources of uncertainty, especially with regards to the presence of clouds that leads to failed images, e.g. up to 50% of images acquired by operational satellites can be trashed due to excessive cloud covers [11]. On this point, a challenge is to be able to exploit both short-term meteorological forecast to prefer dispatching solutions where a given sub-area is allocated to a mission disposing of a satellite that overflies that area over a sunny period, and historical weather data to identify regions whose observation opportunities should not be missed. For this, one ambition is to learn a model of the *long-term reward* provided by a dispatching decision. Another

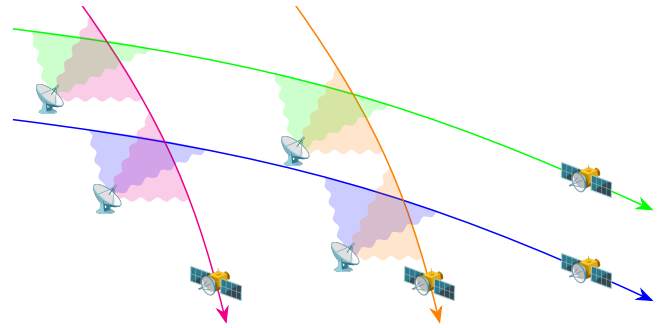


Figure 4: Four satellites, each one having potential contacts with two ground stations

key mechanism is to build a federation layer that iteratively performs task re-dispatching, to automatically update the coverage strategy based on the actual execution status.

2.4 Related works

In the literature, several contributions deal with the observation of large areas [2, 14, 15, 17], but there is a gap between these contributions and the federation layer looked for (e.g., need to manage multiple requests, need to manage the current load of the missions, or need to deal with the cloud coverage). From a larger perspective, defining the multi-mission coverage service is related to several basic research questions such as “how to build a surrogate model of the capacities of a set of agents among which a set of goals must be dispatched?”, “how to estimate the long-term reward associated with a given task allocation given various sources of uncertainty?”, or “how to exploit an *hybrid AI* approach combining optimization and machine learning for the federation of a set of resources?”.

3 FEDERATED COMMUNICATION

For the communication part, we particularly elaborate on the use of non proprietary communication stations using the GSaaS paradigm [3, 16], where customers book communication resources from a ground segment provider, on as a pay-as-you-go basis, instead of building their own stations. We address here the following decision problem: “*How to assign bundles of communication windows to satellites in order to meet data flow requirements, to minimize jamming, and to minimize costs induced by booking services?*” Such a decision has to be made on a regular basis for several days in advance or when an urgent request requires downloading data very quickly.

3.1 System Composition

The problem involves a set of N *satellites* (held by mission agents), seeking for communication windows for data transfer, and a set of M independent communication site agents, referred to as *sites*, implementing GSaaS interfaces.

Within the federation layer, a *communication site booking agent* should be in charge of providing communication opportunities to satellites of federated missions (see Fig. 2). It should manage the communications between N satellites and M sites including ground stations. This management must be based on the concept of *contact*, i.e. communication between a satellite and a site for a

time interval, and this agent must: (1) compute all the potential contacts (Fig. 4 illustrates the concept of potential contact), (2) select the contacts to be booked and (3) book the selected contacts. For point (1), missions must provide orbits of their satellites and an orbit propagator could compute the position of each satellite along time and record events corresponding to entering or exiting the visibility cone of each station. These computations provide, for each satellite i , L_i potential contacts. For point (3), the booking activity of the agent can be purely reactive and based on pre-defined rules that shall be respected when interacting with each GSaaS provider, based on an agreed service level.

3.2 Multi-Agent Decisions and Problems

The selection of contacts to be booked is driven by the communication needs provided by each EO mission. The booking of selected contacts is performed by making requests to agents associated with each ground station network. The actual booking of a contact is not always successful because the requested site may not have a sufficient capacity. Thus, the booking agent may iterate activities (2) and (3) during negotiation rounds with the site agents.

The selection of contacts could be based on solving an optimization problem whose variables, $x_{i,l} \in \{0, 1\}$, would correspond to the use of contact $l \in \{1, \dots, L_i\}$ by satellite i , for $i \in \{1, \dots, N\}$. For each satellite i , the goal is to satisfy K_i contact needs, where each need k is defined by a required contact duration $D_{i,k}$, a list of candidate sites, and a time window. Each need of each satellite shall be satisfied, leading to $\sum_{i=1}^N K_i$ constraints. Two criteria are relevant for selecting contacts: the total cost, C , and the total level of conflict and jamming, J . The cost is the sum of selected contact costs, i.e. $C = \sum_{i=1}^N \sum_{l=1}^{L_i} c_{i,l} x_{i,l}$. With respect to conflict and jamming, those may occur when time windows of two contacts of two satellites on the same site overlap. It can be written as $J = J_{\text{out}} + J_{\text{in}}$ with $J_{\text{out}} = \sum_{i=1}^N \sum_{l=1}^{L_i} f_{i,l} x_{i,l}$ and $J_{\text{in}} = \sum_{i=1}^{N-1} \sum_{l=1}^{L_i} \sum_{j=i+1}^N \sum_{m=1}^{L_j} b_{i,l,j,m} x_{i,l} x_{j,m}$, where $f_{i,l}$ characterizes conflict and jamming between contact l of satellite i and satellites not belonging to the federation. $b_{i,l,j,m}$ characterizes conflict and jamming between contact l of satellite i and contact m of satellite j . If those contacts do not overlap or if they are not related to the same site $b_{i,l,j,m} = 0$. Otherwise, the value of $b_{i,l,j,m}$ depends on site characteristics. If the site has a single antenna, it is a conflict and $b_{i,l,j,m} = 1$. If the site has multiple antennas jamming may occur when the angle between satellites i and j seen from the site is lower than a critical value. In that case, $b_{i,l,j,m}$ could be computed as the ratio of the duration in which this condition holds to the duration of the union of the two contacts. The cost and jamming criteria could be combined using a lexicographic optimization strategy, where one criterion is optimized first and ties are broken by optimizing the second criterion.

3.3 Challenges for the Federation Layer

Highly combinatorial problems. For a booking agent, the problem of selecting contacts, i.e. the slots to be requested to GSaaS, is highly combinatorial, due to the large number of satellites and contacts, and multi-criteria. On the centralized side, several algorithmic options are possible, such as local search, tree search, and integer or

constraint programming. It could also be considered as a *Multi-Agent Resource Allocation* or *Multi-Agent Planning* problem [6, 21]. Yet, dedicated solution methods have to be devised, to scale up when considering hundreds of sites and tenth of satellites for time horizons longer than a week. Moreover, due to possible interference between the allocated contacts, the booking problem becomes a strongly coupled multi-agent allocation problems, where the utility of a single agent strongly depends on other agents' contact bundles.

Need for a model of the other agents. The booking activity may induce a *learning problem* devoted to learning the probability of a contact request being rejected. This is also linked to learning the $f_{i,l}$ terms based on the set of accepted and rejected contacts. As for observation activities, it requires *learning the behavior and the load of each GSaaS service*, and might be addressed as a *Multi-Agent Reinforcement Learning* (MARL) problem [1]. Indeed, the booking agent could learn the behavior of each GSaaS provider depending on past request outcomes, and could also build a surrogate load model to quickly assess the probabilities of acceptance of each booking request, considering that other clients can emit requests for competing time windows.

Strategic behaviors. The booking activity could be impacted by strategic issues. For instance, one may consider booking as soon as possible to increase the probability of acceptance of the booking request or, on the opposite, booking as late as possible to get a better view of the actual communication needs depending on the satellite usages. Again, reinforcement learning techniques could be helpful to optimize such decisions.

3.4 Related Works

The problem consisting in scheduling downloads from satellites of a constellation to a dedicated ground station network is well covered by the literature [4, 5, 19, 23, 24]. More federated visions have been developed in several countries. For instance, ESTRACK is a global network of ground stations that supports ESA missions [8]. It consists of 13 stations located in Europe, Africa, Australia, and South America. Other countries also have similar networks, such as the United States, China, Russia, Japan and India. However, the problem of optimizing communication needs of several constellations using several ground station networks with different access conditions is a new problem resulting from the development of several competing GSaaS networks such as AWS [20] and KSAT [13].

4 CONCLUSION

We propose a multi-agent approach for orchestrating EO activities. This reduces the need for manual intervention and promotes autonomous decisions for observation scheduling and communication allocation. Still, several challenges remain, due to the combinatorial complexity of the problems and the difficulty of modeling the intricate system with learning algorithms. Further research from the AI and AAMAS community is required to address these challenges.

ACKNOWLEDGMENTS

This work is part of the DOMINO-E project which received funding from the European Union's Horizon Europe Programme for Research and Innovation under Grant Agreement n°101082230.

REFERENCES

- [1] Stefano V. Albrecht, Filippos Christianos, and Lukas Schäfer. 2024. *Multi-Agent Reinforcement Learning: Foundations and Modern Approaches*. MIT Press, Cambridge, MA. <https://www.marl-book.com>
- [2] M. Barkaoui and J. Berger. 2020. A new hybrid genetic algorithm for the collection scheduling problem for a satellite constellation. *Journal of the Operational Research Society* 71, 9 (2020), 1390–1410.
- [3] Elisa Carcaillon and Berylia Bancquart. 2020. Market perspectives of Ground Segment as a Service. In *71st International Astronautical Congress (IAC)*. International Astronautical Federation (IAF), Online, 15.
- [4] Jeremy Castaing. 2014. Scheduling downloads for multi-satellite, multi-ground station missions. In *28th Annual AIAA/USU Conference on Small Satellites*. AIAA/USU, Logan, Utah, 12.
- [5] Hao Chen, Baorong Zhai, Jiangjiang Wu, Chun Du, and Jun Li. 2020. A satellite observation data transmission scheduling algorithm oriented to data topics. *International Journal of Aerospace Engineering* 2020 (2020), 1–16.
- [6] Yann Chevaleyre, Paul Dunne, Endriss Ulle, Lang Jérôme, Lemaître Michel, Nicolas Maudet, Julian Padget, Steve Phelps, Juan Rodriguez-Aguilar, and Paulo Sousa. 2006. Issues in Multiagent Resource Allocation. *Informatica* 30 (01 2006).
- [7] Peter Cramton, Yoav Shoham, and Richard Steinberg. 2006. *Combinatorial Auctions*. The MIT Press, Cambridge, MA.
- [8] Sylvain Damiani, Holger Dreihahn, Jörg Noll, Marc Niézette, and Gian Paolo Calzolari. 2006. Automated Allocation of ESA Ground Station Network Services. In *International Workshop on Planning and Scheduling for Space (IWSPSS'06)*. self-published, Baltimore, MD, USA, 4.
- [9] DOMINO-E Consortium. 2024. *DOMINO-E – Your access to multi-mission Earth observation*. Airbus Defence and Space, CapGemini, Tilde, ONERA, ITTI, Oikoplus, Airbus, GMV. Retrieved February 6, 2024 from <https://domino-e.eu/>
- [10] EUSPA. 2022. *EO and GNSS market report*. Technical Report. European Union Agency for the Space Programme (EUSPA).
- [11] Adrien Hadj-Salah, Rémi Verdier, Clément Caron, Mathieu Picard, and Mikaël Capelle. 2019. Schedule Earth Observation satellites with Deep Reinforcement Learning. arXiv:1911.05696 [cs.LG]
- [12] Sarit Kraus. 2001. *Strategic Negotiation in Multiagent Environments*. MIT Press, Cambridge, MA, USA.
- [13] KSAT. 2023. *Ground Network Services*. KSAT. Retrieved December 2, 2023 from <https://www.ksat.no/ground-network-services/>
- [14] C. Lenzen, M. Dauth, T. Fruth, A. Petrak, and E. Gross. 2021. Planning Area Coverage with Low Priority. In *International Workshop on Planning and Scheduling for Space (IWSPSS)*. self-published, Online, 80–88.
- [15] A. Maillard, S. Chien, and C. Wells. 2021. Planning the Coverage of Solar System Bodies Under Geometric Constraints. *Journal of Aerospace Information Systems* 18, 5 (2021), 289–306.
- [16] Louis Nguyen. 2012. *Ground Stations as a Service (GSaaS) for Near Real-time Direct Broadcast Earth Science Satellite Data*. Technical Report. NASA. https://esto.nasa.gov/forums/estf2021/Presentations/June10/Nguyen_GSON_ESTF2021.pdf
- [17] X. Niu, H.Tang, and L. Wu. 2018. Satellite scheduling of large areal tasks for rapid response to natural disaster using a multi-objective genetic algorithm. *Journal of Disaster Risk Reduction* 28 (2018), 813–825.
- [18] Gauthier Picard, Clément Caron, Jean-Loup Farges, Jonathan Guerra, Cédric Pralet, and Stéphanie Roussel. 2021. Autonomous Agents and Multiagent Systems Challenges in Earth Observation Satellite Constellations. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems (Virtual Event, United Kingdom) (AAMAS '21)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 39–44.
- [19] Marco Schmidt. 2011. *Ground station networks for efficient operation of distributed small satellite systems*. Ph.D. Dissertation. Universität Würzburg.
- [20] Amazon Web Services. 2023. *AWS Ground Station*. Amazon. Retrieved December 2, 2023 from <https://aws.amazon.com/ground-station/>
- [21] Yoav Shoham and Kevin Leyton-Brown. 2008. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press, USA.
- [22] S. Tuli, G. Casale, and N. R. Jennings. 2022. GOSH: Task Scheduling Using Deep Surrogate Models in Fog Computing Environments. *IEEE Transactions on Parallel and Distributed Systems* 33, 11 (nov 2022), 2821–2833. <https://doi.org/10.1109/TPDS.2021.3136672>
- [23] Rafael Vazquez, Federico Perea, and Jorge Galán Vioque. 2014. Resolution of an Antenna–Satellite assignment problem by means of Integer Linear Programming. *Aerospace Science and Technology* 39 (2014), 567–574.
- [24] Andrea J Witt. 2018. *Optimization of CubeSat Ground Stations for Increased Satellite Numbers*. Ph.D. Dissertation. Monterey, CA; Naval Postgraduate School.