Avoiding Self-Interference in Megaconstellations through Cooperative Satellite Routing and Frequency Assignment

Nils Pachler, Edward F. Crawley, Bruce G. Cameron

***Abstract*—With the reduced distance between satellites in modern megaconstellations, the potential for self-interference has emerged as a critical challenge that demands strategic solutions from satellite operators. The goal of this paper is to propose a cooperative framework that combines the Satellite Routing (i.e., mapping of beams to satellites) and Frequency Assignment (i.e., mapping of frequency spectrum to beams) strategies to mitigate self-interference both within and between satellites. This approach stands in contrast to current practices found in the literature, which address each problem independently and solely focus on intra-satellite interference. This study presents a novel methodology for addressing the Satellite Routing problem, specifically tailored for modern constellations to maximize capacity while effectively mitigating self-interference through the use of Integer Optimization. By combining this method with established Frequency Assignment techniques, the results demonstrate an increase in throughput of up to 138% for constellations such as SpaceX Starlink. Notably, the study reveals that relying on individual approaches to tackle interference may lead to undesired outcomes, underscoring the advantages of a cooperative framework. Through simulations, the study highlights the practicality and applicability of the proposed method under realistic operational conditions.**

***Index Terms*—Satellite Communications, Resource Management, Integer Programming, Routing, Frequency**

I. Introduction

The increasing prominence of satellite communications is revolutionizing the accessibility and availability of broadband Internet, fundamentally transforming the possibilities of when and how societies can connect. By virtue of their capacity to traverse vast distances without relying on physical infrastructure, satellites present viable solutions in scenarios where ground-based connectivity is lacking or inefficient, such as remote or mobile environments [1]. In recent times, a surge of proposals has emerged, aiming to extend broadband access from space [2], [3]. Notably, prominent companies like SpaceX and Amazon have submitted applications to deploy extensive satellite constellations, pushing the limits of capacity within the satellite industry to unprecedented levels [4]. These new systems represent a tangible and practical alternative to bridging the digital divide and fostering socio-economic development by extending Internet connectivity to traditionally underserved regions [1].

The increased capacity of satellite systems stems from two primary sources. Firstly, advancements in satellite payload technology, including phased array antennas, onboard processing, and routing capabilities, allow for increased utilization of limited spacecraft resources. Secondly, the reduction in launching and manufacturing costs enables larger space segments, significantly enhancing the overall system capacity. Consequently, operators can deploy extensive networks comprising highly capable satellites, effectively addressing the limitations associated with terrestrial infrastructure [5]. However, in this new paradigm where satellites are positioned in close proximity, avoiding self-interference, i.e., interference produced by satellites in the same network, becomes an indispensable prerequisite for seamless operations, demanding utmost attention and strategic mitigation strategies.

To tackle the challenge of self-interference, recent studies have incorporated interference considerations into the resource allocation process of satellite constellations [6], [7]. However, these early investigations suffer from two primary limitations: 1) They mitigate interference within individual satellites rather than between satellites, and 2) They address interference in resources separately (e.g., power, spectrum), only one at a time, without considering the importance of coordinated approaches. Given the intricate interconnected nature of constellation resources, we contend that cooperative strategies encompassing multiple resources are vital for effectively addressing self-interference and maximizing overall performance.

This study aims to bridge the existing research gap by exploring cooperative mechanisms that effectively harness the capabilities of multiple satellites within megaconstellations, while simultaneously mitigating self-interference both within and between satellites. Specifically, our investigation focuses on the application of Satellite Routing, which involves mapping beams to satellites, and Frequency Allocation, which entails allocating spectrum to each beam. The overarching goal is to maximize system performance while ensuring that the interference between signals originating from the same or different satellites remains below a predefined threshold. To accomplish this, we introduce a novel methodology inspired by graph clustering techniques and Integer Optimization to address the Satellite Routing problem effectively. Additionally, by leveraging Frequency Assignment techniques from exist-

Manuscript received MONTH DAY, 2023; revised MONTH DAY, YEAR
Corresponding author: Nils Pachler (pachler@mit.edu)
All authors are with the Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, 02139 Cambridge MA, USA (e-mail: pachler@mit.edu, crawley@mit.edu, bcameron@mit.edu)

TABLE I
SUMMARY OF THE ORBIT CHARACTERISTICS OF THE EXISTING CONSTELLATIONS. VALUES EXTRACTED FROM PUBLIC FILINGS AND MAY BE ALTERED AT THE DISCRETION OF THE OPERATOR. TABLE ADAPTED FROM [5] WITH PERMISSION OF THE AUTHORS.

System	Altitude (km)	Inclination (°)	Planes	Sat. per plane	Num. of sat.	Ref.
O3b mPower	8,062	0	1	10	10	[3]
ViaSat LEO	1,300	45	8	36	288	[8]
Telesat Lightspeed	1,015	98.98	27	13	1,671	[9]
	1,325	50.88	40	33		
SpaceX Starlink	540	53.2	72	22	4,408	[2]
	550	53	72	22		
	560	97.6	6	58		
	560	97.6	4	43		
	570	70	36	20		
OneWeb	1,200	87.9	36	49	6,372	[10]
	1,200	55	32	72		
	1,200	40	32	72		

ing literature, we demonstrate that adopting a collaborative approach yields up to a 75% increase in throughput compared to strategies that consider individual resources in isolation. Furthermore, through extensive simulations, we illustrate the practical viability and effectiveness of our proposed methodology during operational scenarios.

The structure of this paper is as follows: Section II provides a comprehensive literature review, identifying the specific research gap that this study aims to address; Section III outlines the proposed interference model and cooperative framework; Section IV delves into the Satellite Routing formulation, specifically tailored to tackle the challenges of self-interference; Section V presents the methodology developed to effectively address the Satellite Routing problem within realistic operational conditions; Section VI briefly discusses the implementation of the Frequency Assignment problem; Section VII presents the obtained results from applying our framework, emphasizing aspects such as convergence, validation, and performance considerations; Section VIII discusses the operational considerations, crucial when implementing the proposed framework in practical scenarios; finally, Section IX, draws conclusions summarizing the key findings of this study.

II. BACKGROUND AND RELATED WORK

In recent years, the design of satellite constellations has undergone a significant transformation, transitioning from a small number of spacecraft in geostationary orbit (GEO) or medium Earth orbit (MEO) to the deployment of thousands of satellites operating at various altitudes. The characteristics of the O3b mPower,ViaSat LEO, Telesat Lightspeed, SpaceX Starlink, and OneWeb constellation designs are summarized in Table I. These developments have two notable implications in terms of operations, as depicted in Figure 1: 1) Users have multiple satellites within their line-of-sight at all times, enabling operators to select from a set of available satellites to serve each user, and 2) The inclusion of signal interference in the decision-making process is critical to prevent service disruptions due to the close proximity of satellites in large constellations.

TABLE II
ADVANTAGES AND DISADVANTAGES OF APPLYING INTERFERENCE MITIGATION MECHANISMS AT DIFFERENT TIMES

Interference mitigation	Advantages	Disadvantages
Before operations	Reduces real-time control	Higher offline comp. time
	Larger impact	Requires conservative interference estimation
During operations	Addresses only real interference	Higher online comp. time Smaller impact

Signal interference can occur when two closely positioned signals occupy the same frequency spectrum and utilize the same polarization, resulting in information loss and service disruptions. Addressing signal interference involves considering two distinct time-horizons: pre-operations and during operations. In the pre-operations phase, interference can be mitigated by incorporating interference considerations into crucial decisions made prior to the commencement of operations, such as mapping beams to satellites [5] (referred to as the Satellite Routing problem) and assigning frequency spectrum to beams [11] (Frequency Assignment problem). It is important to note that although these problems are typically addressed prior to operations due to their complexity, the solution to these problems forms an operations plan that specifies the assignment of satellites and frequency channels to each beam, which may evolve over time. During operations, interference can be managed through the utilization of interference-aware decision-making mechanisms to handle the power for each beam [12] (Power Allocation problem), by controlling the activation time of each beam [7] (Beam Hopping problem), or by controlling the activation time or frequency spectrum for each user within the beam [13], [14] (multi-frequency time-division multiple-access, MF-TDMA, and Carrier Allocation problems, respectively).

The advantages and disadvantages of implementing interference mitigation at different times are outlined in Table II. Incorporating interference considerations before operations entails increased computational time in the offline phase due to added complexity and necessitates a conservative estimation of interference during actual operations. However, this approach can help minimize the need for real-time control, resulting in enhanced service quality and reduced service dropouts. From an operators perspective, addressing interference before operations, which is the primary focus of this work, holds significant value as it can yield superior system performance and service quality at minimal cost. The following section provides a summary of current approaches employed to address interference in the pre-operations phase.

Previous research has primarily focused on addressing the Satellite Routing and Frequency Assignment problems as separate entities. Specifically, in terms of the Satellite Routing problem, existing literature tends to concentrate on task-based networks characterized by short communication duration, such as phone-based communications [15], [16] and imaging satellites [17]. However, there is a scarcity of research concern-

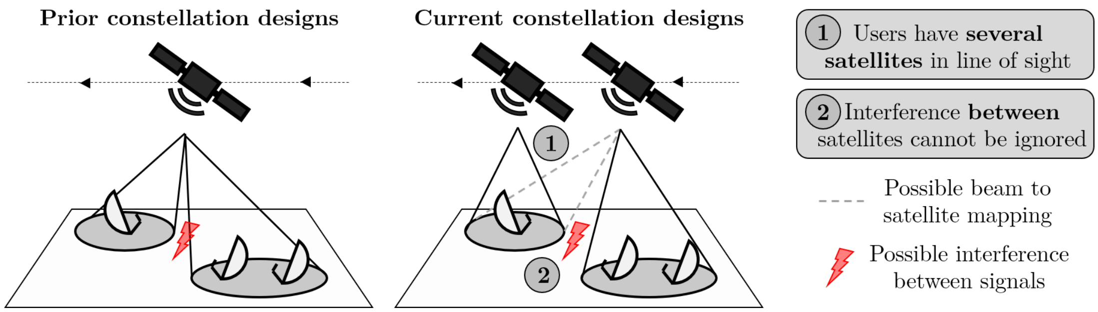

Fig. 1. Implications of current constellation designs

ing broadband satellite communications, which necessitate establishing connections of indefinite duration. In particular, Dai [18], [19] proposes a greedy approach where each user autonomously selects the optimal set of satellites. From the satellite perspective, Yin [20] develops a Deep Reinforcement Learning (DRL) approach to learn the policy for connecting to user stations. To deal with high dimensionality, Jiang [21] addresses the challenge by employing a multi-objective iterative sub-gradient method to simultaneously maximize coverage and capacity. However, these approaches rely on conservative Frequency Assignments to mitigate interference, irrespective of the user-to-satellite mapping, resulting in a reduced system capacity. In contrast, Pachler [22] incorporates intra-satellite interference as an objective to minimize within the Satellite Routing formulation, enabling a more efficient utilization of the frequency spectrum. However, due to the implicit assumptions of the formulation, the proposed method is only applicable to single-plane equatorial constellations.

Regarding the Frequency Assignment problem, existing approaches concentrate on maximizing capacity and mitigating intra-satellite interference within small constellations. These state-of-the-art techniques encompass a wide range of methods, including Neural Networks [23], [24], Deep Reinforcement Learning [25], and Integer Linear Programming [6]. While these approaches offer effective solutions for the specific problem, their suitability for megaconstellations is limited due to the absence of inter-satellite interference consideration. To address this limitation, some studies have introduced a time-independent interference set, providing a workaround for inter-satellite interference. Building upon this concept, Pachler [11] and Garau-Luis [26] present formulations and efficient solution methods explicitly tailored for multi-satellite constellations. However, it is important to note that this interference set is derived from the Satellite Routing policy, which is not within the scope of these works. Consequently, an ongoing challenge lies in developing efficient approaches to obtain an interference set that accurately accounts for the necessary interference when considering the Satellite Routing policy.

The primary objective of this study is to employ collaborative Satellite Routing and Frequency Assignment strategies to maximize capacity while effectively mitigating interference within satellite constellations. As tends to be the case in satellite communications [11], [13], [27], we contend that collaborative strategies, as opposed to individual approaches, are essential for avoiding overly conservative solutions that result in suboptimal resource utilization. This paper contributes to the field in the following ways:

- A novel cooperative framework to mitigate interference and optimize the performance of satellite constellations by combining Satellite Routing and Frequency Assignment strategies.
- A unique methodology to address the Satellite Routing problem within multi-plane satellite constellations with the objective of maximize capacity and minimize interference based on mixed integer linear programming.
- Validation and performance analyses on the SpaceX Starlink and O3b mPower constellations to verify and evaluate the impact of the proposed methodology to address self-interference under realistic operational conditions.
- A comprehensive analysis on the implications of incorporating the presented formulation into operations.

III. Proposed Cooperative Framework

This Section presents the cooperative framework proposed in this work for capacity maximization and self-interference mitigation in non-geostationary orbit (NGSO) constellations. To achieve this objective, we describe the interference model, define the Satellite Routing and Frequency Assignment problems, and outline our cooperative strategy.

A. Interference model

As emphasized in the preceding Section, this study focuses on mitigating interference by including signal considerations into pre-operations decision-making. However, determining whether two signals will interfere depends on various factors such as power usage, atmospheric conditions, and other real-time considerations. Therefore, during the pre-operations phase, it is not possible to definitively ascertain if two signals will interfere. Nonetheless, there are specific conditions that guarantee **non-interference** during operations (see Figure 2): 1) non-overlapping frequency spectrum, 2) utilization of orthogonal polarizations, or 3) sufficient geographic separation. We consider two beams to have *possible interference* if none of these conditions are met. For the purposes of this work, if at least one of these conditions is fulfilled, we conclude that two signals will not interfere or that the interference is sufficiently low to be easily managed during operations. Following this,

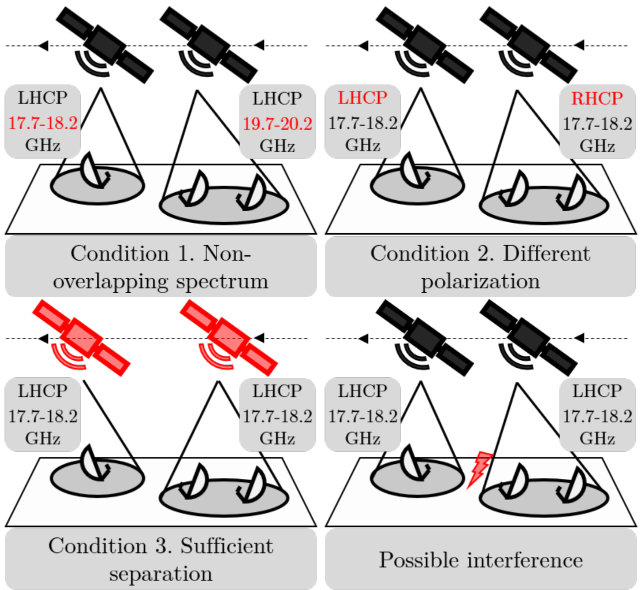

Fig. 2. Conditions to mitigate possible interference. LHCP/RHCP: Left/right hand circular polarization.

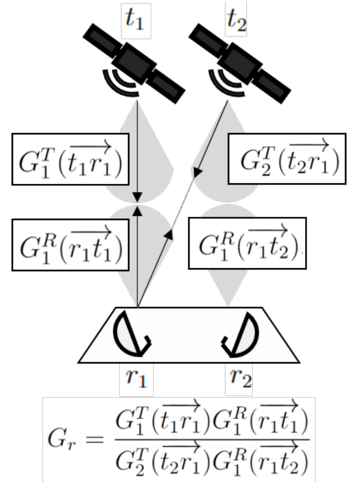

Fig. 3. Definition of relative gain, which applies the concept of isolation, used to evaluate when two beams have sufficient separation (condition 3).

interference is modelled as a binary parameter, representing if two signals potentially interfere or not.

To determine possible interference between two beams, it is necessary to evaluate each of the aforementioned conditions. The first two conditions are determined exclusively by the Frequency Assignment problem: overlapping frequency spectrum is assessed by comparing central frequency and bandwidth, while concurring polarization is determined through a binary comparison. Regarding the third constraint, assessing the proximity of two signals relies on factors such as the physical position of the transmitters and receivers, which are related to the Satellite Routing problem, and the radiation pattern of the transmitting and receiving antennas, which is influenced by the shape of each beam and the specific hardware employed. For the purposes of this study, circular beam shapes with a predefined and known aperture angle δ are assumed.

Existing research has primarily relied on angular distance [11], [28][1] to estimate when beams are *sufficiently separated*. The prevailing approach considers potential interference between two beams if their angular distance falls below a specific threshold, while distances exceeding the threshold are assumed to be interference-free. However, this method solely accounts for intra-satellite interference and does not address inter-satellite interference. To overcome this limitation, we propose an alternative approach that utilizes the concept of isolation (see Figure 3). Instead of relying solely on angular distance, we evaluate the impact of one signal (S_2) on another signal (S_1) by calculating the relative gain of S_2 at the reception point of S_1. The relative gain represents the ratio of the observed gain at the reception of S_1 in the direction of S_2's transmission. If the relative gain exceeds a predefined threshold (I_{thres}), it indicates insufficient separation between the signals and potential interference, while values below the threshold indicate negligible interference. It is important to note that the concept of relative gain is equivalent to angular distance when both signals have the same source.

[1] Commonly used frequency coloring techniques are a simplified application of angular separation

Formally, let the receiver of signal S_1 be located at r_1 with an antenna gain of $G_1^R(p)$ as a function of the direction p. Similarly, we define the position t_1 and the antenna gain function $G_1^T(p)$ for the transmitter of S_1. We apply a similar definition to S_2 by introducing r_2, $G_2^R(p)$, t_2, and $G_2^T(p)$. The gain of signal S_1 at r_1, denoted as G_1, is computed as the product of $G_1^T(\overrightarrow{t_1r_1})$ and $G_1^R(\overrightarrow{r_1t_1})$. Notably, $\overrightarrow{t_1r_1}$ corresponds to the main direction of the gain, resulting in G_1 being the maximum gain. Similarly, we calculate the gain of signal S_2 at r_1, denoted as G_2, using $G_2^T(\overrightarrow{t_2r_1})$ and $G_1^R(\overrightarrow{r_1t_2})$. It is important to note that $\overrightarrow{t_2r_1}$ does not correspond to the main direction of the gain. The relative gain, denoted as G_r, is computed as the ratio of G_1 to G_2. The model for the signal gain, both on- and off-axis, is described in Appendix C. The threshold I_{thres} serves two main purposes: I_{thres} represents the maximum allowable difference in signal strength, computed as $\frac{C}{N+I}$, for the two signals to be considered comparable at reception, and the margin of relative gain that allows for non-negligible co-channel and cross-polarization effects.

B. Relation between the Satellite Routing and Frequency Assignment

The interference mitigation framework proposed in this study adopts a collaborative approach, integrating the Satellite Routing and Frequency Assignment problems. Individually, the Satellite Routing problem is defined as finding the optimal mapping between beams and satellites, considering that each beam may serve multiple users. Conversely, the Frequency Assignment problem is defined as determining the appropriate central frequency and bandwidth allocation for each beam. This decision-making process must also account for system-specific considerations such as frequency reuse and the potential for dual polarization. The Frequency Assignment method must conform to the technical specifications of the system,

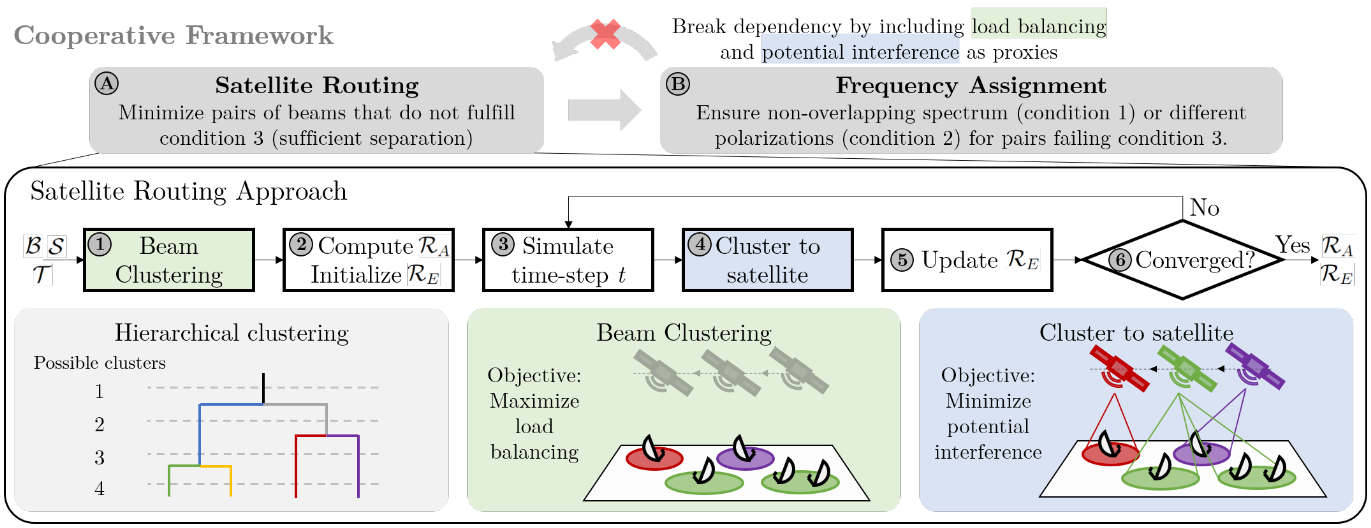

Fig. 4. Visual representation of the proposed cooperative framework and Satellite Routing approach for self-interference mitigation.

ensuring compliance with hardware limitations such as the maximum allowable frequency reuse.

Using these definitions, we can discern two significant interconnections: 1) The Frequency Assignment hinges upon the Satellite Routing solution in order to determine the available resources, and 2) Both problems impact the potential interference between beams.

C. Defining the cooperative framework

Considering the computational intractability of jointly solving these NP-hard problems [22], [29], the proposed approach is to address each problem individually, first tackling Satellite Routing and then addressing Frequency Assignment. Note that this sequencing responds to the fact that the Frequency Assignment exhibits stronger dependencies on the Satellite Routing than vice versa. This approach breaks the dependencies by leveraging two key observations. Firstly, the expected load of each satellite can serve as a proxy for evaluating the quality of the Satellite Routing solution. By favoring solutions that distribute the load more evenly across satellites, efficient resource utilization during the Frequency Assignment phase can be achieved. And secondly, while both problems influence the potential interference, they affect different conditions independently. Notably, if one condition is met, the values of the other two conditions become irrelevant as interference will not occur. Based on this, the proposed framework comprises two main components (refer to Figure 4):

- A Satellite Routing: The quality of the solution is evaluated based on both the load distribution across satellites and the potential interference between beams. A good solution aims to distribute beams across satellites while maximizing sufficient geographical separation between them.
- B Frequency Assignment: Building upon the Satellite Routing solution, this phase incorporates the information obtained and considers potential interference as a critical factor to minimize. A satisfactory solution effectively utilizes available resources and prioritizes non-overlapping spectrum or different polarizations for beams lacking sufficient geographical separation.

Subsequent sections will define the formulation and implementation for each problem. It is important to note that the approach, which addresses each problem individually, aligns with prior literature. However, by adopting a system perspective and analyzing both problems concurrently, we are able to incorporate considerations in the problem formulations that may not be evident when examining them in isolation.

IV. Satellite Routing Formulation

This Section describes the formulation for the Satellite Routing problem, which aims to distribute the workload across satellites and mitigate potential interference. We define the problem using an integer linear formulation and analyze the computational complexity associated with this formulation.

A. Problem definition

The Satellite Routing problem involves determining the mapping between beams and satellites at each time. As discussed in the previous Section, an optimal solution for this problem aims to achieve load balancing among satellites while minimizing potential interference between beams. It is assumed in this study that the user-to-beam mapping guarantees user coverage within the beams footprint, irrespective of the satellites position, as long as it is above a predefined minimum elevation angle. Therefore, meeting this requirement falls outside the scope of the Satellite Routing problem.

Formally, we denote $\mathcal{B}$ as the set of beams that require mapping to a set of satellites $\mathcal{S}$ over a specified duration $\mathcal{T}$. Each beam is characterized by its center p_b and an associated expected demand d_b. Optionally, $\mathcal{B}$ can include information specifying a subset of satellites $\mathcal{S}_b$ that are responsible for covering a particular beam b. If this subset is defined, only satellites within $\mathcal{S}_b$ are eligible for covering beam b. In the absence of this specification, we assume that $\mathcal{S}_b = \mathcal{S}$. The set $\mathcal{S}$ provides information regarding the orbital characteristics of each satellite, with the assumption that orbital elements, except for the mean anomaly, remain fixed during operations.

B. Problem formulation

To tackle the Satellite Routing problem, we introduce the binary variable $x_{b,s}(t)$, representing the mapping between beam b and satellite s at time t. We enforce the constraint that each beam can only be mapped to a single satellite at any given time: $C1 : \sum_{s\in\mathcal{S}_b} x_{b,s}(t) = 1$. Additionally, we ensure that the mapping is restricted to valid and visible satellites for each beam: $C2 : x_{b,s}(t) \leq \mathbb{1}_{s\in\mathcal{S}_b}\mathbb{1}_{s\in LoS(b,t)}$, where $LoS(b,t)$ denotes the set of points within the line of sight (LoS) of beam b at time t. To balance beams across satellites, we introduce the auxiliary binary variable $y^A_{b_1,b_2}$, indicating whether beam b_1 and beam b_2 are active on the same satellite at any point during $\mathcal{T}$. Based on the values of $x_{b,s}(t)$, $y^A_{b_1,b_2}$ can be expressed as:

$$C3 : y^A_{b_1,b_2} \geq x_{b_1,s}(t) + x_{b_2,s}(t) - 1 \quad (1)$$

To address interference between beams, we introduce another auxiliary binary variable, $y^E_{b_1,b_2}$, indicating whether beam b_1 and beam b_2 lack sufficient geographical separation at any point during $\mathcal{T}$. This can be determined using the binary parameter $I_{b_1,s_1,b_2,s_2}(t)$, which indicates if the relative gain between b_1 connected to s_1 and b_2 connected to s_2 exceeds I_{thres}, as explained in Section III-A. Using $x_{b,s}(t)$ and $I_{b_1,s_1,b_2,s_2}(t)$, we can define $y^E_{b_1,b_2}$ as:

$$C4 : y^E_{b_1,b_2} \geq I_{b_1,s_1,b_2,s_2}(t)(x_{b_1,s_1}(t) + x_{b_2,s_2}(t) - 1) \quad (2)$$

To optimize the problem, we introduce weighting factors ω_A and ω_E to balance the load and minimize interference, respectively. The complete formulation is as follows:

$$\begin{aligned} \min_{x_{b,s}(t)} \quad & \sum_{b_1\in\mathcal{B},b_2\in\mathcal{B}} (\omega_A y^A_{b_1,b_2} + \omega_E y^E_{b_1,b_2}) \\ s.t. \quad & C1;\ C2;\ C3;\ C4;\ x_{b,s}(t), y^A_{b_1,b_2}, y^E_{b_1,b_2} \in \{0,1\} \end{aligned} \quad (3)$$

The presented formulation represents an integer-linear formulation of the Satellite Routing problem in NGSO constellations, with the objective to achieve load balancing across satellites while minimizing potential interference. Once a solution is obtained, we define $\mathcal{R}_A$ and $\mathcal{R}_E$ as the sets comprising pairs of beams that coincide on the same satellite or lack sufficient geographical separation, respectively. $\mathcal{R}_A$ is denoted as the *overlapping* set, while $\mathcal{R}_E$ is denoted as the *interference* set. These sets can be derived from the variables $y^A_{b_1,b_2}$ and $y^E_{b_1,b_2}$, and will be utilized during the Frequency Assignment approach.

C. Complexity analysis

To assess the complexity of the proposed formulation, we introduce the parameter M, which represents the maximum number of visible and valid satellites for a single beam at any given time. M is computed as $M = \max_{b,t}\sum_s \mathbb{1}_{s\in\mathcal{S}_b}\mathbb{1}_{s\in LoS(b,t)}$. The memory requirement for the problem is characterized as $\mathcal{O}(|\mathcal{B}|M|\mathcal{T}| + |\mathcal{B}|^2)$, where the first term accounts for the decision variables ($x_{b,s}(t)$) and the second term represents the auxiliary variables ($y^A_{b_1,b_2}$ and $y^E_{b_1,b_2}$). The total number of possible solutions is on the order of $O(M^{|\mathcal{B}||\mathcal{T}|})$, as each beam and time requires selecting one (1) variable from a maximum of M options. Additionally, Appendix A provides a proof demonstrating that the Satellite Routing decision problem is NP-complete.

V. Proposed approach for the Satellite Routing

This Section describes the proposed methodology for tackling the Satellite Routing problem, as outlined in Section IV. Given the inherent computational complexity of achieving an optimal solution, we divide the problem into two subproblems, which enables us to obtain an optimized solution within a reasonable timeframe.

A. Reducing complexity by decomposition

Given the computational complexity of the problem formulation described in Section IV-C, an off-the-shelf mathematical solver becomes infeasible for real-world operational scenarios with tens of thousands of beams, thousands of satellites, and solutions that need to be valid for indefinite duration. To address this challenge, we propose a two-layer decomposition approach that trades optimality for computational efficiency. First, we address the time duration issue by solving the problem incrementally, one time-step at a time. Second, we employ clustering techniques to group beams and assign these clusters to satellites, reducing the combinatorial complexity. The subsequent subsections elaborate on the implementation details of each step.

B. Solving for an indefinite time-horizon

One of the principal complexities of the prior formulation is the exponential growth of possible solutions with the time horizon, which can be indefinite. To mitigate this challenge, we propose a time-step-based approach where the problem is solved iteratively, focusing on one (1) time-step at a time. However, this introduces the challenge of varying values for $y^A_{b_1,b_2}$ and $y^E_{b_1,b_2}$ between different time-steps, which affects the overlapping and interference sets. To address this, we propose an iterative method that aims to identify the smallest sets $\mathcal{R}_A$ and $\mathcal{R}_E$ without requiring a solution for all time-steps. This is achieved by modifying the constraints in Equations 1 and 2 to consider only the pairs of beams not already belonging to $\mathcal{R}_A$ and $\mathcal{R}_E$ during the computation:

$$C3' : y^A_{b_1,b_2} \geq x_{b_1,s}(t) + x_{b_2,s}(t) - 1 \quad \forall \underline{\{b_1,b_2\} \notin \mathcal{R}_A} \quad (4)$$

$$C4' : y^E_{b_1,b_2} \geq I_{b_1,s_1,b_2,s_2}(t)(x_{b_1,s_1}(t) + x_{b_2,s_2}(t) - 1) \quad \forall \underline{\{b_1,b_2\} \notin \mathcal{R}_E} \quad (5)$$

C. Solving a single time-step

The second challenge of the original formulation lies in the significant size of the space segment, particularly for modern constellations where the number of visible satellites per beam can range from tens to hundreds [5]. Prior research has demonstrated the computational intractability of addressing the Satellite Routing directly through integer programming, even for smaller constellations like mPower [22]. Even when addressing a single time-step, the extensive search space and NP-Hardness may hinder the efficient discovery of high-quality solutions. To mitigate this complexity, we introduce the concept of clusters. A cluster denotes a group of beams that share the resources of a subset of satellites. Beams within the same cluster may be mapped to the same group of satellites, while beams in different clusters may not. The main notion behind clustering is to partition the resource pool –originally comprising the entire constellation– into multiple pools (clusters). The original Satellite Routing problem is then decomposed into two sub-problems: beam clustering and cluster-to-satellite mapping. By assuming a low dimensionality of clusters, we convert a high-dimensional to high-dimensional mapping problem into two low-dimensional to high-dimensional mapping sub-problems, thereby reducing complexity. Importantly, Satellite Routing focuses on load distribution across beams, satellites, and clusters, and does **not** aim to optimize resource allocation within individual clusters. The latter falls under the scope of Frequency Assignment.

Formally, a cluster is a group of beams for which any pair of beams within the cluster that shares at least one valid and visible satellite at some point in time is in $\mathcal{R}_A$. Each beam is associated with a single cluster, with the number of available clusters per beam determined by the minimum number of visible and valid satellites across all time steps. To capture nearby beams with varying cluster availability, a hierarchy is established among the clusters, as depicted in the lower left corner of Figure 4. Consequently, the constraint of a single cluster is extended to all clusters within the same branch of the hierarchy. For instance, a beam in the blue cluster forms pairs in $\mathcal{R}_A$ with beams in the black, green, or yellow clusters if they share overlapping satellites, but not with beams in the grey, red, or purple clusters. The set $\mathcal{R}_B$ denotes pairs c_1, c_2 belonging to the same branch, implying potential inclusion in $\mathcal{R}_A$. When assigning clusters to satellites, a beam can only connect to satellites assigned to the same cluster or clusters downstream within the same branch. Hence, a beam in the blue cluster can solely associate with a satellite assigned to the blue, green, or yellow cluster, but not the black, grey, red, or purple clusters. This arrangement guarantees that only beams adhering to the appropriate constraints are assigned to the same satellite. The set $\mathcal{R}_D$ comprises pairs c_1, c_2 where the clusters are on the same branch, with c_2 either equal to or further downstream than c_1.

D. Beam clustering

To assign clusters to beams, we introduce binary variable $x_{b,c}$, representing whether beam b is assigned to cluster c. Each beam can only be assigned to a subset of clusters $\mathcal{C}_b$, contingent upon the minimum number of visible and valid satellites for that beam at any given time. A beam must be allocated to exactly one valid cluster, as expressed by constraint $C5 : \sum_{c\in\mathcal{C}_b} x_{b,c} = 1$. The determination of beam pairs belonging to $\mathcal{R}_A$ is achieved by verifying if they share a satellite that is both valid and visible for both beams simultaneously at any time:

$$\begin{aligned} z_{b_1,b_2} &= \max_{s,t} \mathbb{1}_{s\in LoS(b_1,t)} \mathbb{1}_{s\in\mathcal{S}_{b_1}} \mathbb{1}_{s\in LoS(b_2,t)} \mathbb{1}_{s\in\mathcal{S}_{b_2}} \\ C6 : y^A_{b_1,b_2} &\geq z_{b_1,b_2} \mathbb{1}_{\{c_1,c_2\}\in\mathcal{R}_B} (x_{b_1,c_1} + x_{b_2,c_2} - 1) \end{aligned} \tag{6}$$

It is important to note that the variable z_{b_1,b_2} is solely dependent on the geometric characteristics of the problem and can be computed in advance. This equation allows us to determine which beam pairs will share the same satellite at some point, without explicitly specifying the beam-to-satellite mapping. However, computing the potential interference requires precise knowledge of the beam-to-satellite mapping, as it is highly sensitive to these assignments. Nevertheless, we can utilize the distance between beams to estimate when two beams might be assigned to the same or neighboring satellites. To incorporate this information, we introduce the distance factor f_{b_1,b_2}, defined as follows:

$$\begin{aligned} a_{b_1} &= \max_{b_2} ||p_{b_1} - p_{b_2}|| z_{b_1,b_2} \\ f_{b_1,b_2} &= \left(1 - \frac{||p_{b_1} - p_{b_2}||}{a_{b_1}}\right) z_{b_1,b_2} \end{aligned} \tag{7}$$

By definition, if $p_{b_1} = p_{b_2}$, $f_{b_1,b_2} = 1$, and if $z_{b_1,b_2} = 0$, $f_{b_1,b_2} = 0$. The computation of f_{b_1,b_2} can be incorporated into the geometric analysis of the problem. With this, the objective is to distribute the demand of nearby beams among different clusters, achieving load balancing across satellites while minimizing potential interference. The complete formulation for the beam clustering problem is presented below:

$$\begin{aligned} \min_{x_{b,c}} \quad & \sum_{b_1\in\mathcal{B}, b_2\in\mathcal{B}} (f_{b_1,b_2} y^A_{b_1,b_2}) \\ s.t. \quad & C5;\ C6;\ x_{b,c}, y^A_{b_1,b_2} \in \{0,1\} \end{aligned} \tag{8}$$

Note that, unlike the time-dependent nature of the original formulation (Equation 3), the beam clustering formulation is not time-dependent, requiring only a single solution. By utilizing this formulation, the complete overlapping set $\mathcal{R}_A$ can be computed without the need to simulate all time-steps.

E. Cluster-to-satellite mapping

Once the beam clustering is established, the subsequent step involves the time-dependent cluster-to-satellite mapping. To address this efficiently, we can solve for a single time-step, as previously explained in Section V-B. For a given time t, we introduce the binary variable $x_{s,c}$, representing the assignment of satellite s to cluster c. Ensuring that each satellite is assigned to only one cluster can be formulated as constraint $C7 : \sum_{c\in\mathcal{C}} x_{s,c} = 1$.

In each time-step, every beam maintains a list of potential satellite connections, **sequenced by preference** as $\mathcal{S}^*_b = s_1, s_2, ..., s_k$, where the beam would ideally prefer connecting to s_1 over s_2, s_2 over s_3, and so on. This ordering parameter, customizable by the operator, could, for instance, be based on visible satellites sorted by highest elevation angle to lowest. However, for a beam b with cluster c_b and satellite s with cluster c_s, the beam can only be assigned to satellite s if it has a matching cluster ($\{c_b, c_s\} \in \mathcal{R}_D$). Notably, a beam b will be assigned to satellite s if it shares a matching cluster and if no previous satellite in $\mathcal{S}_b^*$ has such a cluster. This assignment can be efficiently modeled using the binary variable x_{b,s_i}:

$$C8 : x_{b,s_i} = \sum_{\{c_b,c\}\in\mathcal{R}_D} (x_{s_i,c} - \sum_{j=1}^{i-1} x_{s_j,c}) \tag{9}$$

Note that in the aforementioned equation, x_{b,s_i} might assume values less than 1, but we will consider beam b matched with satellite s_i only when $x_{b,s_i} = 1$. Furthermore, each beam must have at least one valid and visible satellite with a matching cluster: $C9 : \sum_{\{c_b,c\}\in\mathcal{R}_D} \sum_{s\in\mathcal{S}_b^*} x_{s,c} \geq 1 \ \forall b$. Utilizing the variable $x_{b,s}$, we can formulate the cluster-to-satellite mapping sub-problem as:

$$\begin{aligned} \min_{x_{s,c}} \quad & \sum_{b_1\in\mathcal{B},b_2\in\mathcal{B}} y^E_{b_1,b_2} \\ s.t. \quad & C4';\ C7;\ C8;\ C9;\ x_{s,c}, y^E_{b_1,b_2} \in \{0,1\} \end{aligned} \tag{10}$$

This refined formulation allows us to address the cluster-to-satellite mapping sub-problem as an integer-linear problem, similar to the original formulation. Notably, the variables and constraints concerning load balancing ($y^A_{b_1,b_2}$, $C4'$) do not appear, as they were already accounted for during the beam clustering phase. Integrating the outcomes of both the beam clustering and cluster-to-satellite mapping phases and the satellite preferences provides a comprehensive solution for the Satellite Routing problem. Nonetheless, in certain scenarios, the feasibility of the prior formulation may be affected by geometric properties of the problem. For methods to tackle such situations, refer to Appendix B.

F. Complete approach

The proposed approach involves the following sequential steps (see Algorithm 1):

1) Assign each beam to a cluster by solving the beam clustering problem in Equation 8.
2) Compute $\mathcal{R}_A$, initialize $\mathcal{R}_E = \emptyset$, and set $t = t_0$.
3) Simulate time t, calculating geometrical parameters.
4) Assign each satellite to a cluster by solving the cluster-to-satellite mapping in Equation 10.
5) Add any new pairs of beams with $y^E_{b_1,b_2} = 1$ to $\mathcal{R}_E$.
6) Upon convergence, finish, otherwise set t to the next time-step and return to step 3.

We check for convergence by verifying if there have been N_{conv} iterations without adding new beam pairs to $\mathcal{R}_A$ or $\mathcal{R}_E$. For the resolution of the formulations (step 1 and 4), we employ commercial mathematical solvers. Although the proposed approach addresses one time-step at a time, $\mathcal{R}_A$ and $\mathcal{R}_E$ capture the dynamic nature of the constellation by tackling a time-invariant problem and simulating multiple time-steps until convergence, respectively. Notably, we exclude the weights in Equation 3 in subsequent formulations. This choice is based on each of the two steps (beam clustering and cluster-to-satellite mapping) focusing on distinct metrics—balancing the load and avoiding interference, respectively—making a weighted approach unnecessary.

Algorithm 1 Satellite Routing approach

Input: $\mathcal{B}, \mathcal{S}$ ▷ Sets of beams and satellites
Input: I_{thres} ▷ Interference threshold
Input: N_{conv} ▷ Iterations without change until convergence
Output: $\mathcal{R}_A, \mathcal{R}_E$ ▷ Overlapping and interference sets
$\mathcal{R}_A$ = Solve eq. 8 ▷ Compute overlapping set through beam clustering
$\mathcal{R}_E = \emptyset$ ▷ Initialize the interference set
$t = t_0$ ▷ Initialize the solution
iterNoImprovement = 0 ▷ Initialize convergence factor
while iterNoImprovement $< N_{conv}$ **do** ▷ While not converged
  Compute LoS $\forall b, s$ ▷ Line of sight computation
  Compute I_{b_1,s_1,b_2,s_2} $\forall b_1, s_1, b_2, s_2$ based on I_{thres} ▷ Interference computation
  $\mathcal{R}'_E$ = Solve eq. 10 ▷ Compute the new interference set
  if $\mathcal{R}'_E == \emptyset$ **then**
    iterNoImprovement = iterNoImprovement + 1 ▷ Increase counter
  else
    iterNoImprovement = 0
  $\mathcal{R}_E = \mathcal{R}_E \cup \mathcal{R}'_E$ ▷ Update the interference set

G. Complexity analysis

For complexity analysis, we introduce N_{CT} and N_{CB}, representing the total number of clusters and the maximum number of clusters a single beam can be mapped to, respectively. Additionally, N_T denotes the total iterations until convergence. The memory requirements for *beam clustering* and *cluster-to-satellite mapping* are $O(|\mathcal{B}|N_{CB} + |\mathcal{B}|^2)$ and $O(|\mathcal{S}|N_{CT} + |\mathcal{B}|^2)$, respectively. The search space size scales with $O(N_{CB}^{|\mathcal{B}|})$ and $O(N_{CT}^{|\mathcal{S}|})$, respectively. We can demonstrate that both formulations remain NP-completeby using a similar reasoning shown in Appendix A. However, our approach significantly reduces complexity due to the reduction in options per beam (with N_{CB} at most M, often smaller) and by addressing one time-step at a time. In Section VII, we confirm that commercial mathematical solvers efficiently solve the Satellite Routing problem using this approach.

VI. Frequency Assignment implementation

To tackle the Frequency Assignment problem, we adopt the method proposed in [26]. In that study, the authors present an integer-linear formulation that assigns central frequency, bandwidth, and reuse factor to the beams, considering constraints derived from $\mathcal{R}_A$ and $\mathcal{R}_E$ as defined in this work. Note that the frequency assignment is assumed to remain fixed for an indefinite duration. Although it could theoretically be possible to adjust the spectrum each time a beam changes satellite, potentially enhancing efficiency, the complexity of addressing such a problem for high-dimensional scenarios is currently intractable, with uncertain benefits. The objective of this formulation is to minimize a set of objectives, such as power consumption or bandwidth, while mitigating potential interference between the beams, aligning perfectly with the objectives of this study. We will utilize this implementation without modifications, focusing specifically on minimizing power consumption. A concise overview of the Frequency Assignment approach follows and it is summarized in Algorithm 2.

The objective of the Frequency Assignment is to assign the initial frequency, bandwidth, and reuse factor for each

Algorithm 2 Frequency Assignment approach

Input: $\mathcal{B}$ ▷ Set of beams
Input: $\mathcal{R}_A, \mathcal{R}_E$ ▷ Overlapping and interference sets
Input: N_{ch} ▷ Number of beams allowed to change per iteration
Input: N_{cutoff} ▷ Number of options allowed per beam and bandwidth
Input: N^{FA}_{conv} ▷ Iterations without change until convergence
Output: f_b, w_b, g_b ▷ Frequency, bandwidth, and reuse factor for each beam

$f_b, w_b, g_b \forall b \in \mathcal{B}$ = Warm start [11] ▷ Initialize the solution
iterNoImprovement = 0 ▷ Initialize convergence factor
previousObjective = 0 ▷ Initialize previous objective
while iterNoImprovement $< N_{conv}$ **do** ▷ While not converged
  $b_0 = Rand(\mathcal{B})$ ▷ Select one beam randomly
  $\mathcal{B}' = Neighbors(b_0, N_{ch})$ ▷ Select N_{ch} closest neighbors to i
  $\mathcal{X} = \emptyset$ ▷ Initialize pool of variables
  for $b \in \mathcal{B}'$ **do** ▷ For each non-fixed beam
    Find $l_{f,w,g,b}$ for each valid option
    $vars = RankAndCutOff(l_{f,w,g,b}, N_{cutoff})$ ▷ Restrict options to the most interesting ones
    if $Valid(f_b, w_b, g_b)$ **then** ▷ If current solution is valid
      $vars = vars \cup \{(f_b, w_b, g_b)\}$ ▷ Add it to the pool
    $\mathcal{X} = \mathcal{X} \cup vars$
  Solve Eq. 11 only with variables in $\mathcal{X}$, fixing $\mathcal{B} \setminus \mathcal{B}'$ ▷ Solve the reduced problem
  $f_b, w_b, g_b \forall b \in \mathcal{B}'$ = Solution ▷ Update with the new solution
  if currentObjective == previousObjective **then**
    iterNoImprovement = iterNoImprovement + 1 ▷ Increase counter
  else
    iterNoImprovement = 0
  previousObjective = currentObjective

beam. To that end, $x_{f,w,g,b}$ is a binary variable indicating if beam b is mapped to initial frequency f, bandwidth w, and reuse factor g. We define $l_{f,w,g,b}$ as a utility function representing the expected power consumption of beam b when using frequency f, bandwidth w, and reuse factor g, computed as in [26]. We also introduce s_{f_1,w_1,f_2,w_2} as a binary parameter determining if the spectrum bounded by f_1 and w_1 overlaps with f_2 and w_2: $C10 : s_{f_1,w_1,f_2,w_2} = \mathbb{1}_{f_1 \leq f_2+w_2} \mathbb{1}_{f_2 \leq f_1+w_1}$. We enforce constraints to ensure that beams sharing a satellite (i.e., a pair in $\mathcal{R}_A$) do not share overlapping resources: $C11 : s_{f_1,w_1,f_2,w_2}(x_{f_1,w_1,g,b_1} + x_{f_2,w_2,g,b_2}) \leq 1 \forall \{b_1, b_2\} \in \mathcal{R}_A$. Furthermore, beams susceptible to potential interference (i.e., a pair in $\mathcal{R}_E$) must not share overlapping spectrum: $C12 : s_{f_1,w_1,f_2,w_2}(x_{f_1,w_1,g_1,b_1} + x_{f_2,w_2,g_2,b_2}) \leq 1 \forall \{b_1, b_2\} \in \mathcal{R}_E$. Finally, each beam can only be assigned, at most, once: $C13 : \sum_{f,w,g} x_{f,w,g,b} \leq 1$. The problem then reduces to finding the set of $x_{f,w,g,b}$ that maximize the utility function:

$$\begin{aligned} \min_{x_{f,w,g,b}} \quad & \sum_{f,w,g,b} (l_{f,w,g,b} - M_{\mathcal{B}}) x_{f,w,g,b} \\ s.t. \quad & C10; C11; C12; C13; x_{f,w,g,b} \in \{0,1\} \end{aligned} \tag{11}$$

Here, $M_{\mathcal{B}}$ serves as a large number aimed at maximizing the number of active beams, defined as $M_{\mathcal{B}} = \max_{f,w,r,p,b} l_{f,w,r,p,b} |\mathcal{B}|$. To tackle the Frequency Assignment, we initially obtain a warm start using the method in [11]. Subsequently, we iteratively optimize a subset of N_{ch} beams until no superior solution is found for N^{FA}_{conv} steps. During each iteration only the most promising set of options is considered for each beam, using a ranking approach based on the utility function. Combining the Satellite Routing and Frequency Assignment methodologies achieves the desired outcome: Satellite Routing divides beams into resource pools, while Frequency Assignment optimizes the utilization of each pool. Notably, the Frequency Assignment remains agnostic to the cluster-like representation employed in Satellite Routing, as all necessary information is encoded in the overlapping and interference sets.

VII. Performance Results

In this Section, we present the results obtained from the cooperative Satellite Routing and Frequency Assignment framework proposed in this study. We begin by demonstrating the validation and convergence analysis of the Satellite Routing approach. Subsequently, we conduct a performance analysis on five contemporary constellations to evaluate the effectiveness of the proposed cooperative framework and its impact.

A. Experimental set-up

To assess the proposed methodology under realistic operational conditions, we adopt representative space and user segments from modern environments. As this work focuses on fixed satellite service (FSS) environments providing broadband connectivity, we utilize four established constellation designs for the space segment, namely O3b mPower, ViaSat LEO, Telesat Lightspeed, SpaceX Starlink, and OneWeb. To simulate the user segment, we generate a user distribution based on the global population distribution [30]. This involves creating a grid with a resolution of $0.1° \times 0.1°$ and generating N_{loc} locations, where the selection probability of each location corresponds to the percentage of population residing in that cell. At each location, we assume $N_{us/loc}$ users, each requiring a continuous 100 Mbps connection. The specific values of N_{loc} and $N_{us/loc}$ vary depending on the experiment. For the O3b mPower constellation, only users within the +/- 50° latitude range are considered valid. The users are then organized into beams using the algorithm outlined in [11].

Given the growing interest in mobile satellite services (MSS), we also explore the performance of our proposed methodology in such systems. Specifically, we focus on the SpaceX Starlink constellation, which have partnered with T-Mobile for their mobile services. The user generation follows a similar procedure, with each user requiring a 100 Kbps connection. For the purposes of this study, we assume that users assigned to specific beams remain within the footprint of the beam at all times, and that the beam center remains fixed over time. The orbital configuration of each constellation is presented in Table I. The payload configuration for each satellite is outlined in Table III. The parameters for each experiment are summarized in Table IV. In all simulations, I_{thres} and N_{conv} are set to -30dB and 10 iterations, respectively. To address the different formulations, we execute the commercial solver Gurobi [31] (version 9.1.2) in an Intel(R) Xeon(R) Platinum 8160 CPU @ 2.10GHz, allowing up to 16 simultaneous threads. For the purpose of this work, $\mathcal{S}^*_b$ has been simulated as a list of visible satellites ordered by elevation angle, highest to lowest.

TABLE III
SUMMARY OF THE LINK CHARACTERISTICS OF THE VARIOUS CONSTELLATIONS. CHARACTERISTICS MARKED WITH * ARE EXTRACTED FROM PUBLIC FILINGS OR PUBLIC INFORMATION. REST OF THE VALUES ESTIMATED BY THE AUTHORS. EIRP: EFFECTIVE ISOTROPIC RADIATION POWER

Characteristic	O3b mPower	ViaSat LEO	Telesat Lightspeed	SpaceX Starlink	OneWeb	SpaceX Starlink
Service	FSS	FSS	FSS	FSS	FSS	MSS
User Downlink band [GHz]	17.8 - 18.6* 18.8 - 20.2*	17.8 - 18.6* 18.8 - 20.2* 37.5 - 42*	17.8 - 18.6* 18.8 - 19.3* 19.8 - 20.2*	10.7 - 12.7*	10.7 - 12.7*	1.910 - 1.915* 1.990 - 1.995*
Satellite antenna gain [dB]	44.25	52.4*	32.5*	34*	37*	18.4
EIRP density (P_{max}) [dBW/Hz]	-34.3	-43.7*	-50*	-51.1*	-49.4*	-20.9
Frequency reuse factor	10	3	4	4	4	4
User terminal diameter [m]	1*	0.6*	1*	0.6*	1*	0.0511
Pointing loss [dB]	0.6*	0.1*	0.1*	0.1*	0.4*	0.1*
Rotation loss [dB]	0.5*	0.1*	0.0*	0.1*	1.0*	0.1*
Beam aperture angle [°]	2	2	2	2	2	2
Minimum (desired / emergency) elevation angle [°]	25/10*	25/25*	25/10*	25/25*	50/25*	25/25*

TABLE IV
SUMMARY OF THE SIMULATIONS PERFORMED.

Experiment	Constellation	N_{loc}	$N_{us/loc}$	N_{ts}	T_s
Verification & Convergence	Starlink	20,000	1	10	120
Performance	mPower ViaSat Lightspeed Starlink OneWeb	20,000	1 10 10 10 10	10	120
Operations	Starlink	10,000	10	200	5

TABLE V
SUMMARY OF THE IMPLEMENTATIONS USED.

Satellite Routing	Frequency Assignment
Maximum elevation angle	Water filling [11]
Particle Swarm Optimization [22]	Integer Optimization [26]
Integer Optimization (this work)	

We assess the performance of the proposed methods by initially solving the Satellite Routing and Frequency Assignment, followed by simulating each constellations satellite-user links for N_{ts} time-steps, with a time interval of T_s seconds between steps. Note that, while our method can be adapted for the uplink, our evaluations focus on downlink simulations. Results are evaluated through two primary metrics: throughput and power consumption, computed as the sum of the individual values for each satellite across the entire constellation, and averaged over all time-steps. For insights into the metrics of individual beams, refer to Appendix C. As secondary metrics, the number of active satellites represents the average count of satellites with non-zero throughput, and the spectrum represents the sum of the bandwidth assigned for all beams.

B. Convergence and Validation analysis

This experiment serves a dual purpose: firstly, to investigate the validity of the proposed convergence criteria, and secondly, to assess the effectiveness of the Satellite Routing method. For this evaluation, we simulate the SpaceX Starlink constellation with parameters $N_{loc} = 10,000$, $N_{us/loc} = 1$, $N_{ts} = 10$, and $T_s = 120$.

Figure 5 illustrates the evolution of interference set ($\mathcal{R}_E$) and throughput until convergence is achieved. The beam clustering phase, aimed at determining a complete overlapping set ($\mathcal{R}_A$) representing pairs of beams sharing a satellite at some point in time, occurs once, as explained in Section V-F. This phase spans the initial 10,000 seconds. Subsequently, during the cluster-to-satellite phase, we observe the dynamic evolution of the interference set, reflecting potential interference among beam pairs over time. This iterative phase gradually increases the size of the interference set. Towards the final iterations, the growth in $\mathcal{R}_E$ stabilizes, validating the convergence criteria. Convergence is achieved around 120,000 seconds, affirming the effectiveness of the criteria to obtain a solution in reasonable time. Throughout the simulations, intermediate throughput computations assume an incomplete interference set as if it were complete. However, an incomplete interference set implies potential unaccounted interference. To address this, we randomly disconnect one of the beams in each affected pair during simulations, ensuring no interference albeit at a throughput cost. It is observed that throughput stabilizes after approximately 30,000s, as the majority of constraints have been captured. While further iterations only yield minor improvements in throughput, those are crucial to ensure uninterrupted user service.

Figure 6 provides a snapshot of the constellation state over the Iberian peninsula, displaying the elevation angles of each beam and the respective connected satellites. This visualization confirms the feasibility of the proposed solution, as no elevation angle falls below the operator-defined minimum (25°). Furthermore, the distribution of beams across satellites

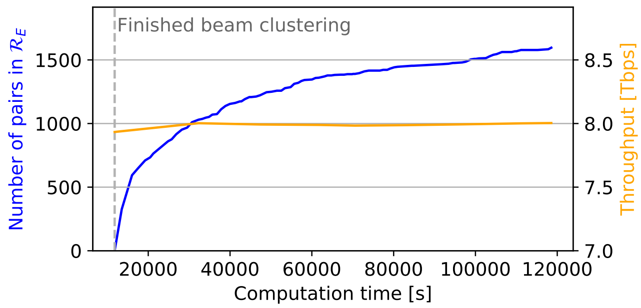

Fig. 5. Convergence analysis on the SpaceX Starlink constellation. The image shows the evolution of the $\mathcal{R}_E$ set (blue) and throughput (orange) over time.

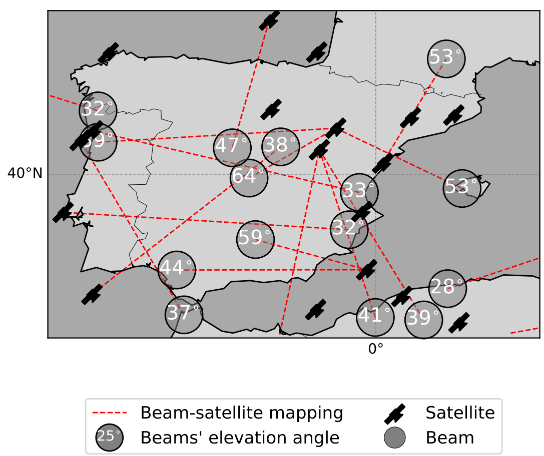

Fig. 6. Validation analysis on the SpaceX Starlink constellation. The image shows one instance of beams, represented with their footprints, mapped to satellites. Each connection is represented through a dashed red line. The value inside the circles represent the elevation angle of the assigned satellite w.r.t. the beam center. Footprints and satellites not at scale.

demonstrates successful load balancing, optimizing resource utilization. Notably, the assignment of nearby beams to different satellites aligns with the objectives of the formulation: achieving load balance across satellites and minimizing potential interference, thereby enhancing overall performance and service quality. Hence, this validation confirms the feasibility of the solution and its alignment with the intended objectives.

C. Performance analysis

The purpose of this experiment is to evaluate the impact of the cooperative methodology compared to individual state-of-the-art approaches derived from research. We assess three different implementations for the Satellite Routing problem: 1) Mapping each beam to the satellite with the highest elevation angle, 2) Utilizing the particle swarm optimization algorithm (PSO) from [22] (only applicable to the O3b mPower constellation), and 3) Our proposed cooperative approach detailed in this work. For the Frequency Assignment problem, we consider two implementations: one heuristic method [11] and one optimized approach [26]. Refer to Table V for a comprehensive list of the implemented methodologies. This comparative analysis aims to shed light on the performance benefits and efficiency of the cooperative approach. This analysis simulates the different constellations with $N_{ts} = 10$ and $T_s = 120$s.

Figure 7 illustrates the simulation results for throughput and power consumption when providing FSS service using the O3b mPower, Telesat Lightspeed, and SpaceX Starlink constellations, along with MSS service using the SpaceX Starlink constellation. Additionally, Table VI provides further details on key performance metrics for each scenario, in addition to results on the ViaSat LEO and OneWeb constellations.

The proposed cooperative framework, represented by the blue dot in Figure 7 and as the IO-IO combination in Table VI, consistently achieves the highest throughput, except for the O3b mPower constellation. This discrepancy is influenced by two factors. Firstly, the low number of visible satellites in LoS at any given point over the Earth surface (typically between 1 and 2 for O3b mPower) limits flexibility during Satellite Routing resolution, resulting in reduced benefits compared to larger systems. Secondly, utilizing a complex Satellite Routing algorithm increases power consumption compared to the maximum elevation angle heuristic, due to greater distances between antennas and consequently higher free space loss, leading to reduced throughput.

Increasing the number of visible satellites enhances flexibility in the Satellite Routing algorithm, compensating for higher path loss and enabling significant throughput gains. Compared to heuristics, the proposed coordinated approach boosts FSS service throughput by approximately 67% in ViaSat LEO, 68% in Telesat Lightspeed, 138% in SpaceX Starlink, and 83% in OneWeb. Notably, throughput improvements are more pronounced with larger constellations, although not directly proportional. However, due to increased path loss, the proposed approach raises power consumption by 143% in ViaSat LEO, 228% in Telesat Lightspeed, 395% in SpaceX Starlink, and 104% in OneWeb.

When employing individual optimization algorithms independently, higher throughput is observed compared to heuristics but is reduced compared to the coordinated method. In systems with more than 1,000 satellites, utilizing only the optimized Satellite Routing (represented by the red dot) provides most of the benefits of the coordinated approach, resulting in throughput increases ranging between 39% and 95%. However, it also increases power consumption by between 106% and 347%. In these large systems, the impact of Frequency Assignment on throughput diminishes, contributing to throughput gains of 14% to 27% but with a comparatively modest increase in power consumption, ranging from 16% to 26%. Conversely, in smaller systems with up to 1,000 satellites where the flexibility of Satellite Routing is constrained, Frequency Assignment assumes a dominant role and becomes the primary driver of throughput improvements.

While counter-intuitive, in medium to large constellations (>100 satellites), adopting the Satellite Routing methodology proposed in this study enables greater spectrum utilization compared to employing the optimized Frequency Assignment method. This outcome is attributed to the ability of the proposed method to distribute demand more effectively, resulting in a higher number of active satellites and consequently increasing the available spectrum pool. With the coordinated

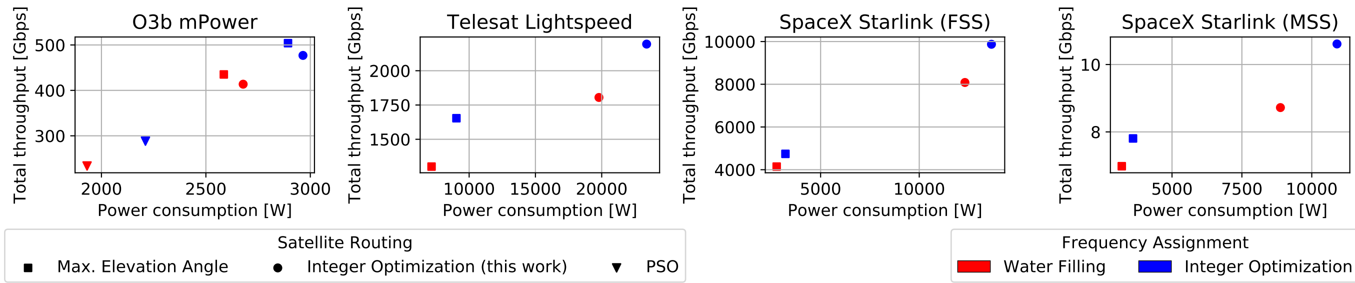

Fig. 7. Throughput and power consumption for different combinations of interference-aware implementations. From left to right, the images show the performance on the O3b mPower, Telesat, SpaceX Starlink when offering FSS service, and SpaceX Starlink when offering MSS service.

TABLE VI
RESULTS FOR VARIOUS FIGURES OF MERIT FOR THE SIMULATIONS OF O3B MPOWER, VIASAT LEO, TELESAT LIGHTSPEED, SPACEX STARLINK, AND ONEWEB. EACH COMBINATION OF SATELLITE ROUTING (SR) AND FREQUENCY ASSIGNMENT (FA) IMPLEMENTATIONS IS REPRESENTED IN A ROW. † REPRESENTS THE SR METHODOLOGY PROPOSED IN THIS WORK. ACRONYMS: ME - MAXIMUM ELEVATION ANGLE; IO - INTEGER OPTIMIZATION; PSO - PARTICLE SWARM OPTIMIZATION; WF - WATER FILLING.

SR	FA	System (Service)	Throughput [Gbps]	Power [kW]	Active Sats.	Spectrum [GHz]	Comp. Time [min]	System (Service)	Throughput [Gbps]	Power [kW]	Active Sats.	Spectrum [GHz]	Comp. Time [min]
ME	WF	O3b mPower (FSS)	435	2.58	10	207	3.1	ViaSat LEO (FSS)	362	**0.086**	93.9	888	**187**
ME	IO		**504**	2.89	10	**246**	12		520	0.104	90.7	1040	197
IO†	WF		414	2.68	10	207	45		402	0.169	**127**	1040	437
IO†	IO		477	2.96	10	231	47		**606**	0.209	**127**	**1270**	440
PSO	WF		234	**1.93**	10	143	**2.5**						
PSO	IO		288	2.21	10	180	5.8						
ME	WF	Telesat Lightspeed (FSS)	1300	**7.13**	449	1590	**645**	SpaceX Starlink (FSS)	4150	**2.75**	937	2170	**1662**
ME	IO		1650	9.00	457	1950	650		4740	3.19	900	2840	1667
IO†	WF		1810	19.8	789	3620	1130		8080	12.3	**1460**	4750	1950
IO†	IO		**2190**	23.4	**807**	**4390**	1137		**9870**	13.6	**1460**	**7260**	1967
ME	WF	OneWeb (FSS)	4660	**6.36**	1230	3410	**237**	SpaceX Starlink (MSS)	6.98	**3.19**	968	8.78	**3783**
ME	IO		5610	7.38	1210	4760	245		7.81	3.59	944	10.6	3883
IO†	WF		7800	13.1	1700	6070	258		8.72	8.87	**1520**	15.4	13350
IO†	IO		**8530**	13.0	**1720**	**8640**	265		**10.6**	10.9	1500	**20.6**	13583

approach, spectrum utilization can be enhanced by up to a factor of 3.3 when compared to heuristic techniques.

In the context of the SpaceX Starlink constellation providing MSS, we note that the throughput increase is comparatively lower at 52% when compared to FSS at 138%. This difference stems from the fact that MSS utilize less directive antennas, leading to reduced effectiveness in frequency reuse. Consequently, the utilization of optimized algorithms for Satellite Routing and Frequency Assignment results in a diminished benefit. Nonetheless, employing the coordinated approach still yields a notable throughput increase. Note that, due to heightened interference among signals, computation time for MSS services extends to weeks due to the need for more simulations to capture the networks dynamism accurately. Conversely, for FSS services, heuristics and optimized methods provide results within minutes, hours, or days, depending on constellation size. Additionally, it is noteworthy that the Particle Swarm Optimization (PSO) technique introduced in [22] is surpassed in performance by the cooperative framework presented in this work. This reinforces the effectiveness and superiority of the proposed approach, demonstrating its potential for real-world implementation.

VIII. OPERATIONAL IMPLEMENTATION

This Section discusses the practical implementation of the proposed methodology from the perspective of a satellite operator. Specifically, it addresses the necessary computations and telemetry required to ensure efficient and feasible operations.

A. Centralized computation

The proposed cooperative framework comprises two main components:

1) Satellite Routing: addresses the allocation of beams to satellites for each time-step. In this approach, operators are required to compute the solution in a centralized manner iteratively, considering all satellites simultaneously. It is crucial to ensure that the time required for each iteration is shorter than the interval between iterations to maintain smooth operations.

2) Frequency Assignment: determines the central frequency, bandwidth, polarization, and frequency reuse for each beam. Unlike the Satellite Routing component, the Frequency Assignment solution remains fixed over time. As a result, operators can compute this information in a centralized manner once, prior to the start of operations.

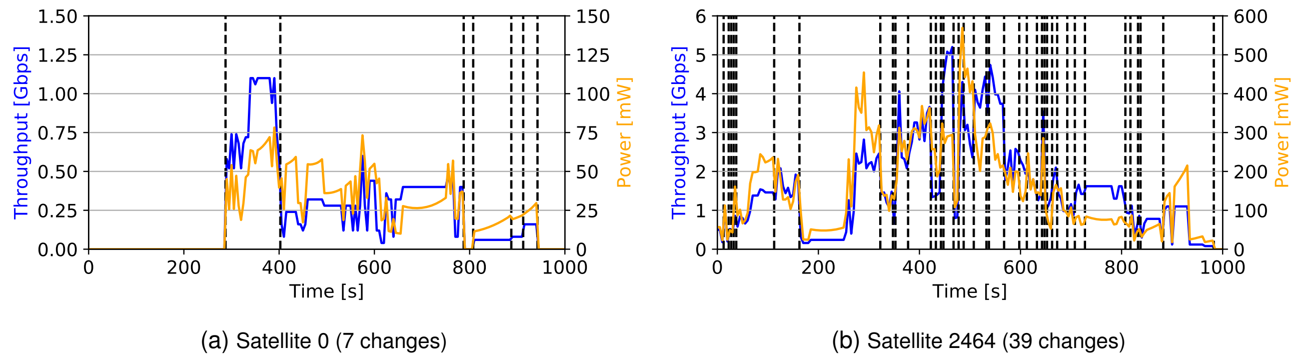

Fig. 8. Throughput and power consumption for two satellites during operations. The dashed lines represent changes in the cluster-to-satellite mapping. The left image corresponds to the first satellite, while the right image corresponds to the satellite with the maximum number of changes.

TABLE VII
SUMMARY OF PARAMETERS AND RESULTS IN THE OPERATIONS SIMULATION

Simulation Parameters		User Parameters		Results			
N_{ts}	200	N_{loc}	10,000	Total number of changes	20,114	Total number of handovers	213,069
T_s (Total simulated time)	5s (1,000s)	$N_{us/loc}$	10	Total computation time	950s	Total number of handovers using maximum elevation	147,471

B. Telemetry updates

Upon computing the solution, the relevant information must be transmitted to the satellites.

1) Satellite Routing: Three blocks of information are involved. First, the cluster assigned to each beam, which remains constant over time and can be transmitted once before operations. Second, the list of satellites that each beam can connect to at each time-step, denoted as $\mathcal{S}_b^*$. If $\mathcal{S}_b^*$ is structured as a complete list of satellites ordered by elevation angle, it can be computed through onboard computation. Alternative preference structures may require additional telemetry. Third, the cluster assigned to each satellite, which is subject to change over time and requires regular updates. Nevertheless, operators do not need to transmit the full mapping of satellites to clusters at every time-step. Instead, they can transmit only the sequence of clusters for each satellite and the times when the cluster changes. This approach optimizes telemetry usage as consecutive time-steps usually have similar solutions. Additionally, the formulation can be optimized to minimize changes at each iteration, further reducing the necessary telemetry (see Appendix D for details).

2) Frequency Assignment: The spectrum information for each beam can be transmitted once before operations and remains fixed throughout the operational phase.

C. Operations simulation

We conducted simulations on the SpaceX Starlink constellation with $N_{loc} = 10,000$, $N_{us/loc} = 10$, $N_{ts} = 200$, and $T_s = 5$s. We analyzed the telemetry requirements and computation time necessary for the proposed approach.

Figure 8 illustrates the evolution of throughput and power consumption for two distinct satellites in the constellation, along with the time placement of cluster-to-satellite changes. As depicted, a cluster change leads to significant fluctuations in both throughput and power consumption. Notably, the fluctuations in periods of no cluster change (e.g., 400s-800s in the left plot), resulting from the dynamic nature of the constellation, changing the beam-satellite visibility conditions, are similar in magnitude as the fluctuations on cluster change. The reason for this is that, under the conditions considered, the number of beams per satellite is relatively low, leading to a high impact in power and throughput when individual beams connect and disconnect from the satellite. Table VII presents the specific results of the simulation. Notably, we observed a total of 20,114 cluster changes during the entire simulation (1,000s). This translates to an average of approximately 20 changes per second or an average time of 219 seconds before each satellite changes clusters. For the satellite with the highest number of changes, a cluster change occurred roughly every 25.6 seconds on average.

Assuming clusters can be encoded using 2 Bytes of information and the time at which the cluster changes can be encoded using 8 Bytes, this implies a total of 1,609 bps to transmit the changes to the whole constellation. Comparing the total number of handovers with simpler methods like the maximum elevation angle heuristic, we found that they are on the same order of magnitude, meaning that the proposed methodology does not introduce unfeasible overhead to the satellites. Importantly, the computation time (950s) required is lower than the total simulated time (1,000s), indicating that continuous operations can be maintained without interruptions. Note that, to ensure feasibility, we do not require the computation on each individual time-step to be lower than the time between time-steps, as the computation can be performed in advance. Continuous operations can be maintained as long as the average time-step computation is shorter than the time between time-steps. These findings demonstrate the feasibility and practicality of the proposed methodology for real-world applications, ensuring efficient use of resources and maintaining optimized performance over time.

IX. Conclusion

This paper presents a pioneering methodology to address the challenges of self-interference in megaconstellations by introducing a comprehensive cooperative framework that encompasses both Satellite Routing and Frequency Assignment decisions. In contrast to existing literature that focuses solely on individual problems and intra-satellite interference, this research leverages the concept of isolation to address interference within and between satellites by solving these resource allocation problems in an integrated manner. Introducing a novel formulation and methodology for the Satellite Routing problem utilizing Integer optimization, we partition the original problem into two clustering problems that can be efficiently solved using state-of-the-art mathematical solvers. By leveraging existing Frequency Assignment implementations from the literature, we demonstrate that our proposed cooperative framework consistently outperforms alternative approaches in terms of total throughput for constellations with more than 100 satellites. Additionally, we validate that the telemetry requirements and real-time computational demands associated with our methodology are feasible, ensuring uninterrupted operations. The conclusions of this work can be summarized as follows:

- The proposed cooperative framework consistently achieves remarkable throughput gains, surpassing other approaches for constellations with more than 100 satellites. Notably, it delivers between 67% and 138% capacity increase in FSS for constellations like ViaSat LEO, Telesat Lightspeed, SpaceX Starlink, and OneWeb.
- Relying solely on individual optimized approaches for Satellite Routing or Frequency Assignment can lead to unfavorable outcomes, such as increased power consumption and reduced throughput.
- To maximize capacity, optimizing Satellite Routing becomes crucial for large constellations (>100 satellites), while optimizing Frequency Assignment becomes crucial for small constellations ($\leq$ 100 satellites).
- While the main gain of the proposed approach is observed in FSS, MSS also benefits from this methodology, with up to 52% throughput increase.
- Despite requiring additional computation and telemetry, the proposed methodology maintains practicality with low resource overhead, ensuring continuous operations.

In light of these findings, the proposed methodology demonstrates its efficacy in effectively harnessing the advantages offered by satellite constellations, thereby addressing self-interference, enhancing satellite utilization, and promoting augmented capacity in modern satellite constellation designs.

Acknowledgments

This work was supported by SES. The authors would like to thank SES for their input to this paper and their financial support. The project that gave rise to these results also received the support of a fellowship from "la Caixa" Foundation (ID 100010434). The fellowship code is LCF/BQ/EU21/11890113. The Satellite Routing clustering approach has been submitted for patent application to the US Patent and Trademark Office, number US 63/631,587.

Author Contributions

Nils Pachler formalized the idea for the cooperative framework, developed the formulation for the Satellite Routing problem, and ran all the analyses. All authors reviewed the manuscript and contributed to its final version.

Appendix A
Proof of NP-completeness

The NP-completeness of the Satellite Routing decision problem can be demonstrated by first showing that the problem is a generalization of the maximum independent set problem, indicating NP-hardness. Subsequently, we establish the existence of a polynomial-time verifier, confirming that the problem is in NP.

To prove NP-hardness, we utilize an oracle that provides the complete sets $\mathcal{R}_A$ and $\mathcal{R}_E$ for the optimal solution. We then define each potential beam-to-satellite mapping ($x_{b,s}(t)$) as a vertex in a graph. Edges are introduced between two nodes $x_{b_1,s_1}(t_1)$ and $x_{b_2,s_2}(t_2)$ if: 1) They address the same beam at the same time (i.e., $b_1 = b_2$ and $t_1 = t_2$), 2) They address the same satellite at the same time, and the pair is not in $\mathcal{R}_A$ (i.e., $s_1 = s_2$, $t_1 = t_2$, and $\{b_1, b_2\} \notin \mathcal{R}_A$), or 3) They present potential interference at the same time, and the pair is not in $\mathcal{R}_E$ (i.e., $I_{b_1,s_1,b_2,s_2}(t_1) = 1$, $t_1 = t_2$, and $\{b_1, b_2\} \notin \mathcal{R}_E$). Solving the Satellite Routing problem is then equivalent to finding an independent set of size $|\mathcal{B}||\mathcal{T}|$ in this graph. Since all nodes related to beam b at time t form a clique, this is equivalent to finding the maximum independent set in the graph, which is NP-complete [32]. As $\mathcal{R}_A$ and $\mathcal{R}_E$ are also variables of the problem, we have established the NP-hardness of the Satellite Routing problem.

To prove NP-completeness, we outline a polynomial-time verifier that, given $x_{b,s}(t)\ \forall b, s, t$, verifies if $\omega_A|\mathcal{R}_A| + \omega_E|\mathcal{R}_E| \leq k$ for $k \in \mathbb{R}_+$. The verifier accomplishes this by traversing the graphs edges and constructing $\mathcal{R}_A$ and $\mathcal{R}_E$ based on $x_{b,s}(t)$, a task achievable in linear time with respect to the number of edges. This proof establishes the Satellite Routing decision problem as being in NP, conclusively demonstrating its NP-completeness.

Appendix B
Ensuring valid clusters

In some instances, the formulation presented in Equation 10 may not yield a feasible solution. This can occur when the mapping between beams and clusters is overly optimistic, resulting in more assigned clusters than available ones. To handle such cases, we introduce an auxiliary binary variable, denoted as x_b, which takes the value 1 if a beam **cannot** find a valid and visible satellite with a matching cluster, and 0 otherwise. We calculate x_b using the constraint $C10$: $\sum_{\{c_b,c\}\in\mathcal{R}_D} \sum_{s\in\mathcal{S}_b^*} x_{s,c} + x_b \geq 1$. Then, we solve:

$$\begin{aligned} \min_{x_{s,c}} \quad & \sum_{b \in \mathcal{B}} x_b \\ s.t. \quad & C7;\; C8;\; C10;\; x_{s,c}, x_b \in \{0,1\} \end{aligned} \tag{12}$$

For beams that could not be matched with a valid satellite, we apply a *downgrading* of the beams cluster assignment. In essence, we reassign the cluster of the beam to the one immediately above in the hierarchy. We then iteratively solve Equation 12 and continue the downgrading process for beams until all beams can be successfully matched with a valid satellite. The resulting solution for $x_{s,c}$ serves as a warm start for Equation 10. Notably, since we are reassigning clusters, $\mathcal{R}_A$ must be recomputed after the downgrading process. This enhanced approach ensures that feasible solutions are always attainable. As a note, in our experiments, around 20% of the beams needed downgrading at some point during the simulation.

Appendix C
Signal model

To model the power usage, we assume that satellites employ adaptive modulation and coding techniques, allowing for the utilization of various modulation and coding (MODCOD) schemes. To compute the power for a beam b, we first determine the appropriate MODCOD based on the spectral efficiency (Γ_b) necessary to fulfill the demand within the available bandwidth, given by $\Gamma_b = \frac{d_b(1+\alpha)}{w_b}$, where α denotes the roll-off factor (set to 0.1). We choose the MODCOD $m_b$whose spectral efficiency is the lowest such that $\Gamma_{m_b} \geq \Gamma_b$. Subsequently, we compute the power for each beam (P_b) using the formula:

$$\begin{aligned} P_b = \frac{E_b}{N}\bigg|_{m_b} \frac{d_b}{w_b} + OBO_{m_b} - G_{TX} - G_{RX} + L_{FSPL} + \\ L + 10\log_{10}(kT_{sys}) + \gamma \end{aligned} \tag{13}$$

Here, we denote $\frac{E_b}{N_0}\big|_{m_b}$ as the energy per bit to noise spectral density, and OBO_{m_b} as the output back-off power of the MODCOD. G_{TX} and G_{RX} represent the gain of the transmitting and receiving antennas, respectively. L_{FSPL} signifies the free-space path loss, and T_{sys} stands for the system temperature, calculated using the Friis equation. Furthermore, γ represents the margin (set to 0.5 dB), and L encompasses various loss components, including atmospheric losses computed through ITU-R P.618 [33], constellation-dependent pointing and rotation losses, and constellation-independent waveguide (0.2 dB) and feeding (1.1 dB) losses.

We assume that each beam has a maximum power consumption proportional to the assigned bandwidth, referred to as the EIRP density (P_{max} in dBW/Hz). The satellites are assumed capable of providing maximum power per beam for all associated beams simultaneously if needed. By not modeling a satellite-wise power limit, we simplify the Power Allocation process, reducing it to solving the link budget equation presented above. If the power required to close the link budget equation P_b exceeds the maximum power per beam ($P_b > (P_{max} - G_{max})w_b$), the power is capped at the maximum value, and the MODCOD with the highest spectral efficiency that closes the link is selected. The throughput is then computed as $R_b = \min(d_b, \frac{\Gamma_{m_b} w_b}{1+\alpha})$.

Moreover, due to the formulations in Satellite Routing and Frequency Assignment, we ensure that two beams will never utilize the same hardware resources, and any signal interference remains negligible ($<$ 30 dB). Notably, our framework models interference at the application layer, before signal generation. This approach can potentially complement existing beamforming models like zero-forcing [34], which handle interference at the physical layer, or real-time interference management techniques like MF-TDMA [13], enhancing flexibility and performance. However, integrating these combined approaches is beyond the scope of this paper and is left out as future work.

Concerning antenna gains, we consider two types: parabolic and isotropic. The gain of a parabolic antenna is calculated as follows: $G_{max} = \eta\left(\frac{\pi D f}{c}\right)^2$, where η denotes efficiency (assumed as 0.65), D is the diameter, f is the transmitting frequency, and c is the speed of light. Antenna off-axis gain follows recommended radiation patterns from the International Telecommunication Union (ITU) for circular beams with a fixed aperture angle, as outlined in ITU-R S.465 [35] for Earth station antennas and ITU-R S.1528 [36] for NGSO satellite antennas. On the other hand, the gain of an isotropic antenna remains constant at -1.5 dB in all directions. We assume satellites employ parabolic-like antennas, while FSS users use parabolic antennas and MSS users use isotropic antennas.

Appendix D
Operational formulation

To minimize the number of necessary changes between consecutive operations, we first compute $x_{s,c}^{t-1}$ and $x_{s,c}^{t}$ by solving Equation 10 for time $t-1$ and time t, respectively. Throughout operations, both solutions will satisfy the constraint $\sum_{b_1 \in \mathcal{B}, b_2 \in \mathcal{B}} y_{b_1,b_2}^E = 0$. Then, we introduce a new auxiliary binary variable, denoted as x_s, which takes the value 1 when satellite s changes clusters from time $t-1$ to time t, and 0 otherwise. We calculate x_s using the constraint $C11 : x_s \geq x_{s,c}^{t-1} - x_{s,c}^{t}\ \forall c$. With these considerations, we formulate the problem to minimize the number of changes as follows:

$$\begin{aligned} \min_{x_{s,c}^t} \quad & \sum_{s \in \mathcal{S}} x_s \\ s.t. \quad & C4';\; C7;\; C8;\; C9;\; C11;\; y_{b_1,b_2}^E = 0;\; x_{s,c}^t, x_s \in \{0,1\} \end{aligned} \tag{14}$$

Where $x_{s,c}^{t-1}$ are fixed variables. This approach enables us to significantly reduce the amount of cluster changes between consecutive operations, optimizing the stability and continuity of the satellite constellations performance over time.

References

[1] G. V. Research, "Video streaming market size, share & trends analysis report by streaming type, by solution, by platform, by service, by revenue model, by deployment type, by user and segment forecasts, 2022-2030," Grand View Research, Tech. Rep., 2022.

[2] Space Exploration Holdings, LLC, "SAT-MOD-20230215-00036," http://licensing.fcc.gov/myibfs/forwardtopublictabaction.do?file_number=SATMOD2023021500036, 2023, accessed: 2023-03-10.
[3] O3b Limited, "SAT-MOD-20200526-00058," http://licensing.fcc.gov/myibfs/forwardtopublictabaction.do?file_number=SATMOD2020052600058, 2020, accessed: 2023-03-09.
[4] N. S. Research, "VSAT and Broadband Satellite Markets," Northern Sky Research, Tech. Rep., 2019.
[5] N. Pachler, I. del Portillo, E. F. Crawley, and B. G. Cameron, "An Updated Comparison of Four Low Earth Orbit Satellite Constellation Systems to Provide Global Broadband," in *IEEE International Workshop*, 2021.
[6] L. Houssin, C. Artigues, and E. Corbel, "Frequency allocation problem in a sdma satellite communication system," *Computers & Industrial Engineering*, vol. 61, no. 2, pp. 346–351, 2011.
[7] P. Angeletti, D. Fernandez Prim, and R. Rinaldo, "Beam hopping in multi-beam broadband satellite systems: System performance and payload architecture analysis," in *24th AIAA International Communications Satellite Systems Conference*, 2006, p. 5376.
[8] Viasat Inc., "SAT-MPL-20200526-00056," http://licensing.fcc.gov/myibfs/forwardtopublictabaction.do?file_number=SATMPL2020052600056, 2020, accessed: 2023-03-08.
[9] Telesat Canada, "SAT-MPL-20200526-00053," http://licensing.fcc.gov/myibfs/forwardtopublictabaction.do?file_number=SATMPL2020052600053, 2020, accessed: 2023-03-09.
[10] WorldVu Satellites Limited, "SAT-MPL-20210112-00007," http://licensing.fcc.gov/myibfs/forwardtopublictabaction.do?file_number=SATMPL2021011200007, 2021, accessed: 2023-03-09.
[11] N. Pachler de la Osa, M. Guerster, I. del Portillo Barrios, E. Crawley, and B. Cameron, "Static beam placement and frequency plan algorithms for leo constellations," *International Journal of Satellite Communications and Networking*, vol. 39, no. 1, pp. 65–77, 2021.
[12] Y. S. Nasir and D. Guo, "Multi-agent deep reinforcement learning for dynamic power allocation in wireless networks," *IEEE Journal on Selected Areas in Communications*, vol. 37, no. 10, pp. 2239–2250, 2019.
[13] Y. Kawamoto, T. Kamei, M. Takahashi, N. Kato, A. Miura, and M. Toyoshima, "Flexible resource allocation with inter-beam interference in satellite communication systems with a digital channelizer," *IEEE Transactions on Wireless Communications*, vol. 19, no. 5, pp. 2934–2945, 2020.
[14] C. He, Y. Hu, Y. Chen, and B. Zeng, "Joint power allocation and channel assignment for noma with deep reinforcement learning," *IEEE Journal on Selected Areas in Communications*, vol. 37, no. 10, pp. 2200–2210, 2019.
[15] Y. Lin, W. Bao, W. Yu, and B. Liang, "Optimizing user association and spectrum allocation in hetnets: A utility perspective," *IEEE Journal on Selected Areas in Communications*, vol. 33, no. 6, pp. 1025–1039, 2015.
[16] Z. Wu, F. Jin, J. Luo, Y. Fu, J. Shan, and G. Hu, "A graph-based satellite handover framework for leo satellite communication networks," *IEEE Communications Letters*, vol. 20, no. 8, pp. 1547–1550, 2016.
[17] Y. Wang, M. Sheng, W. Zhuang, S. Zhang, N. Zhang, R. Liu, and J. Li, "Multi-resource coordinate scheduling for earth observation in space information networks," *IEEE Journal on Selected Areas in Communications*, vol. 36, no. 2, pp. 268–279, 2018.
[18] C.-Q. Dai, J. Luo, S. Fu, J. Wu, and Q. Chen, "Dynamic user association for resilient backhauling in satellite–terrestrial integrated networks," *IEEE Systems Journal*, vol. 14, no. 4, pp. 5025–5036, 2020.
[19] C.-Q. Dai, S. Li, J. Wu, and Q. Chen, "Distributed user association with grouping in satellite–terrestrial integrated networks," *IEEE Internet of Things Journal*, vol. 9, no. 12, pp. 10 244–10 256, 2021.
[20] Y. Yin, C. Huang, D.-F. Wu, S. Huang, M. Ashraf, Q. Guo, and L. Zhang, "Deep reinforcement learning-based joint satellite scheduling and resource allocation in satellite-terrestrial integrated networks," *Wireless Communications and Mobile Computing*, vol. 2022, 2022.
[21] H. Jiang, H. Wang, Y. Hu, and J. Wu, "Dynamic user association in scalable ultra-dense leo satellite networks," *IEEE Transactions on Vehicular Technology*, vol. 71, no. 8, pp. 8891–8905, 2022.
[22] N. Pachler, E. F. Crawley, and B. G. Cameron, "Beam-to-satellite scheduling for high throughput satellite constellations using particle swarm optimization," in *IEEE Aerospace Conference*, 2022.
[23] S. Salcedo-Sanz, R. Santiago-Mozos, and C. Bousoño-Calzón, "A hybrid hopfield network-simulated annealing approach for frequency assignment in satellite communications systems," *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, vol. 34, no. 2, pp. 1108–1116, 2004.
[24] S. Salcedo-Sanz and C. Bousoño-Calzón, "A hybrid neural-genetic algorithm for the frequency assignment problem in satellite communications," *Applied Intelligence*, vol. 22, no. 3, pp. 207–217, 2005.
[25] X. Hu, S. Liu, R. Chen, W. Wang, and C. Wang, "A deep reinforcement learning-based framework for dynamic resource allocation in multibeam satellite systems," *IEEE Communications Letters*, vol. 22, no. 8, pp. 1612–1615, 2018.
[26] J. J. Garau-Luis, S. Aliaga, G. Casadesus, N. Pachler, E. Crawley, and B. Cameron, "Frequency plan design for multibeam satellite constellations using linear programming," *IEEE Transactions in Wireless Communications*, 2022.
[27] N. Pachler, J. J. G. Luis, E. F. Crawley, and B. G. Cameron, "A Unified Resource Allocation Framework and Impact Evaluation for NGSO satellite constellations," *Aerospace Science and Technology (Under review)*, 2024.
[28] L. Lei, E. Lagunas, Y. Yuan, M. G. Kibria, S. Chatzinotas, and B. Ottersten, "Beam illumination pattern design in satellite networks: Learning and optimization for efficient beam hopping," *IEEE Access*, vol. 8, pp. 136 655–136 667, 2020.
[29] T. Mizuike and Y. Ito, "Optimization of frequency assignment," *IEEE Transactions on Communications*, vol. 37, no. 10, pp. 1031–1041, 1989.
[30] Center for International Earth Science Information Network, "Gridded Population of the World, Version 4 (GPWv4): Population Count, Revision 11," https://doi.org/10.7927/H4JW8BX5, Palisades, NY, 2018.
[31] Gurobi Optimization, LLC, "Gurobi Optimizer Reference Manual," 2023. [Online]. Available: https://www.gurobi.com
[32] M. R. Garey and D. S. Johnson, ""Strong"np-completeness results: Motivation, examples, and implications," *Journal of the ACM (JACM)*, vol. 25, no. 3, pp. 499–508, 1978.
[33] Recommendation ITU-R P.618 (2023), "Propagation data and prediction methods required for the design of Earth-space telecommunication systems," 2023.
[34] T. Yoo and A. Goldsmith, "On the optimality of multiantenna broadcast scheduling using zero-forcing beamforming," *IEEE Journal on selected areas in communications*, vol. 24, no. 3, pp. 528–541, 2006.
[35] Recommendation ITU-R S.465 (2010), "Reference radiation pattern for earth station antennas in the fixed-satellite service for use in coordination and interference assessment in the frequency range from 2 to 31 GHz," 2010.
[36] Recommendation ITU-R S.1528 (2001), "Satellite antenna radiation patterns for non-geostationary orbit satellite antennas operating in the fixed-satellite service below 30 GHz," 2001.