Interference-Aware Beam Hopping Scheduling for LEO Satellite Networks Using Soft-Conflict Graph Coloring

Jihane Lghoury
Ericsson R&D, France
Université de Toulouse, Toulouse-INP, CNRS, IRIT, France
jihane.lghoury@ericsson.com

Guilherme Iecker Ricardo
Université de Toulouse, Toulouse-INP, CNRS, IRIT, France
guilherme.ieckerricardo@toulouse-inp.fr

Ioannis Xirouchakis
Ericsson R&D, France
ioannis.xirouchakis@ericsson.com

Xiaotian Fu
Ericsson R&D, France
xiaotian.fu@ericsson.com

Sebastian Euler
Ericsson Research, Sweden
sebastian.euler@ericsson.com

Anestis Tsakmalis
Ericsson R&D, France
anestis.tsakmalis@ericsson.com

Riadh Dhaou
Université de Toulouse, Toulouse-INP, CNRS, IRIT, France
riadh.dhaou@toulouse-inp.fr

***Abstract*—Non-Terrestrial Networks (NTNs), particularly Low Earth Orbit (LEO) satellite systems, face significant intra- and inter-satellite interference challenges. We propose an adaptive, interference-aware beam scheduling framework that accounts for pairwise interference between beams while reducing temporal resource fragmentation. The problem is formulated as a graph coloring model, where nodes represent beams and edges represent interference relationships. The objective is to minimize beam hopping time slots while ensuring a minimum SINR level. A parameter α controls the set of interferers in the graph to reduce the problem complexity. The proposed framework is evaluated and compared against conventional regular beam hopping as well as greedy heuristics, using multiple performance metrics. Numerical results show comparable performance to fixed beam hopping, with slight user throughput satisfaction gains in some scenarios. Moreover, the approach indicates the number of time slots required to achieve satisfactory SINR levels while avoiding unnecessary temporal resource division.**

***Index Terms*—NTN; LEO satellites; beam scheduling; beam hopping; graph coloring**

I. Introduction

Non-terrestrial networks (NTNs) have gained increasing attention for complementing terrestrial networks (TNs) by providing coverage to hard-to-reach areas, thereby enhancing service availability and network resilience. Low Earth Orbit (LEO) satellite constellations use multibeam antennas, giving rise to both intra- and inter-satellite interference. In this context, Beam Hopping (BH) has emerged as a promising beam scheduling technique, where only a subset of beams is activated in each time slot, enabling better adaptation to non-uniform traffic and lower power consumption [1]. However, several approaches of BH rely on distance-based constraints to mitigate interference between simultaneously active beams [2]. While such approaches are effective in TNs, their applicability to NTNs is limited due to the 3D nature of satellite networks, where interference depends not only on 2D spatial separation but also on altitude, satellite movement and the directional characteristics of antenna radiation patterns. These factors result in heterogeneous and asymmetric interference relationships between beams. Subsequent works investigated BH under different constraints. For example, [3] studies dynamic BH for LEO satellite systems with differentiated traffic demands. The problem maximizes a fairness-oriented metric across cells, while considering interference avoidance. Also, in [4], the BH problem is formulated as a throughput-driven or satisfaction-driven optimization problem for integrated satellite–terrestrial networks. However, in these approaches, although interference is taken into consideration, it does not explicitly structure the beam scheduling decisions. More recently, [5] proposed a BH approach combining graph coloring and Tabu Search. Starting from an initial satellite-to-cell assignment, the method iteratively reduces interference conflicts by relaxing the interference threshold until a conflict-free solution is reached. The dynamic satellite–cell assignment adopted in [5] accounts for time-varying satellite coverage. However, at shorter scheduling timescales, the coverage can be considered quasi-static, allowing the satellite–cell association to be fixed and the focus to shift to beam scheduling. In this work, we formulate beam scheduling in LEO satellite systems as an interference-aware graph coloring problem. Cells are modeled as nodes and interference relationships as weighted edges. Unlike [5], the satellite-to-cell association is assumed fixed over the scheduling period. Our objective is to minimize the number of time slots while keeping the aggregate interference experienced by each cell below a predefined threshold. This approach jointly controls SINR and resource utilization, providing a flexible scheduling framework for NTN systems. The remainder of this paper presents the system model, the proposed graph-coloring formulation, the experimental setup and numerical results, followed by the conclusion.

979-8-3195-0167-7/26/$31.00 ©2026 IEEE

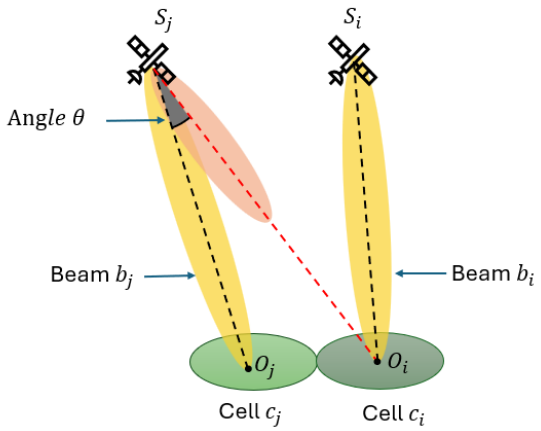


Fig. 1. System model illustrating desired and interfering satellite beams.

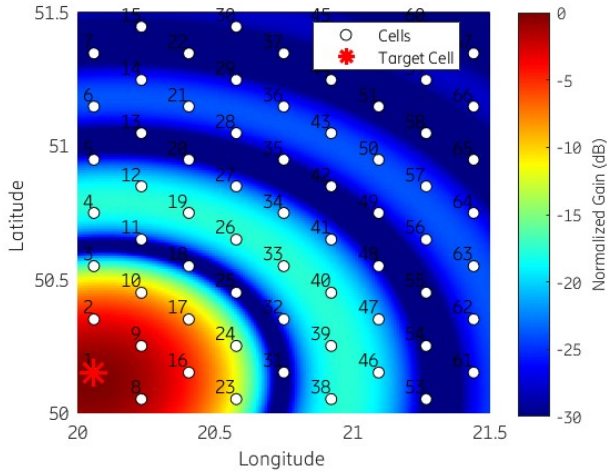


Fig. 2. Beam radiation pattern of a target cell.

II. System Model

Consider a LEO satellite constellation offering a set $\mathcal{B} = \{1, \ldots, B\}$ of beams. We assume that satellites move along predefined and uniform trajectories at constant speed. UEs are uniformly distributed across a service area, which is partitioned into a set $\mathcal{C} = \{1, \ldots, C\}$ of *Quasi-Earth-Fixed* cells of equal size. In 3GPP, *Quasi-Earth-Fixed* cells are fixed with respect to a certain location on the ground for a specific period of time. In this case, satellites generate steerable beams whose footprints remain fixed temporarily on the ground despite the satellite's motion. Cells are planned such that one beam is mapped to each cell at a given time. If beam b is assigned to cell i, its source satellite antenna is directed toward i's center at coordinate O_i, cf. Fig 1.

A. Wireless Propagation Model

Based on Friis' transmission formula [6], the amount of power, in Watts, received at cell i from beam j is modeled as

$$p_{i,j}^{\mathrm{Rx}} = p_j^{\mathrm{Tx}} \cdot g_{i,j}^{Tx} \cdot g_i^{\mathrm{Rx}} \left(\frac{\lambda}{4\pi \| \mathrm{S}_j - \mathrm{O}_i \|} \right)^2, \tag{1}$$

where p_j^{Tx} is the transmission power of beam j, $\mathrm{S}_j \in \mathbb{R}^3$ denotes the satellite position, O_i the center of cell i, λ the carrier wavelength, and g_i^{Rx} the receive antenna gain (assumed identical for all users). The term $g_{i,j}^{Tx}$ represents the transmit antenna gain and depends on the misalignment angle $\theta = \angle \mathrm{O}_i \, \mathrm{S}_j \, \mathrm{O}_j$, i.e., the angle between the beam boresight (direction $\mathrm{S}_j \rightarrow \mathrm{O}_j$) and the direction toward cell i. As specified in 3GPP [7], the antenna gain for a reflector with circular aperture is given by

$$g_\theta = \begin{cases} 1, & \theta = 0°, \\ 4 \cdot \left| \dfrac{J_1(k_0 \cdot a \cdot \sin\theta)}{k_0 \cdot a \cdot \sin\theta} \right|^2, & 0° < |\theta| \leq 90°, \end{cases} \tag{2}$$

where:

- $J_1(\cdot)$ denotes the Bessel function of the first kind and first order;
- a is the radius of the antenna circular aperture;
- $k_0 = \frac{2\pi f}{c}$ is the wave number, where f is the carrier frequency and c is the speed of light.

Figure 2 illustrates the resulting beam gain radiation pattern: While most of the radiated power is concentrated in the main lobe toward the target cell, non-negligible power is also distributed in other spatial directions, causing interference. Note that in (1), if $i = j$, then $\theta = 0°$ and $p_{i,i}^{\mathrm{Rx}}$ corresponds to the desired received power at cell i. We summarize all received power contributions in the matrix $\mathbf{p}^{\mathrm{Rx}} = \{p_{i,j}^{\mathrm{Rx}}, \forall i, j \in \mathcal{C}\}$.

B. Physical Interference Model

We say that cell i is illuminated by its assigned beam if transmission is effectively taking place. The SINR of cell i with respect to a set $\mathcal{C}' \subseteq \mathcal{C}$ of illuminated cells is

$$\sigma_i(\mathcal{C}') = \frac{p_{i,i}^{\mathrm{Rx}}}{\mathcal{N}_0 + \sum_{j \in \mathcal{C}' \setminus \{i\}} p_{i,j}^{\mathrm{Rx}}} \tag{3}$$

where $\mathcal{N}_0$ is the noise floor and $p_{i,j}^{\mathrm{Rx}}$ is defined in (1). Finally, we consider an upper bound for a cell's channel capacity based on Shannon's equation. Then, for a given set of illuminated cells $\mathcal{C}'$, the capacity of cell i is

$$c_i(\mathcal{C}') = \omega \cdot \log_2(1 + \sigma_i(\mathcal{C}')), \tag{4}$$

where ω is the bandwidth of the beam illuminating a cell, assumed to be equal for all beams. The interference-related QoS constraint is enforced through a minimum SIR requirement, equivalently expressed as an interference budget limiting the aggregate interference per scheduled beam.

III. Optimization Problem Definition

Consider that our system operates in a TDMA fashion, such that time is split into κ slots. We must choose in which time slot each cell will be illuminated. As the satellite moves, the set of cells covered by its footprint evolves over time. However, over a time window Ω of a few seconds, this set can be assumed unchanged. The satellite therefore serves a fixed set

of cells during Ω, and the beam scheduling problem is solved over this finite scheduling horizon. Let the *illumination plan* be the set of variables

$$\mathbf{x} \in \{0,1\}^{C\times\kappa} \tag{5}$$

indicating whether cell i is illuminated at time slot k (i.e., $x_{i,k}=1$) or not (i.e., $x_{i,k}=0$). We also define set $\mathcal{C}_k(\mathbf{x}) \subseteq \mathcal{C}$ as the set of cells that are illuminated at time slot k, such that

$$\mathcal{C}_k = \{i \in \mathcal{C} : x_{i,k} = 1\}. \tag{6}$$

Finally, based on (4), we define the system's *overall capacity* as follows:

$$\begin{aligned} c_\kappa(\mathbf{x}) &\triangleq \frac{1}{\kappa}\sum_{k=1}^{\kappa}\sum_{i\in\mathcal{C}} c_i(\mathcal{C}_k(\mathbf{x})) \\ &= \frac{1}{\kappa}\sum_{k=1}^{\kappa}\sum_{i\in\mathcal{C}} w\cdot\log_2(1+\sigma_i(\mathcal{C}_k(\mathbf{x})) \\ &= \frac{w}{\kappa}\sum_{k=1}^{\kappa}\sum_{i\in\mathcal{C}} \log_2\left(1+\frac{p^{\text{Rx}}_{i,i}}{\mathcal{N}_0+\sum_{j\in\mathcal{C}_k(\mathbf{x})\setminus\{i\}} p^{\text{Rx}}_{i,j}}\right) \end{aligned} \tag{7}$$

Problem 1. In this problem, our objective (8) is to find the number of time slots κ^* and the illumination plan $\mathbf{x}^*$ that maximize $c(\mathbf{x})$ such that (i) each cell is illuminated once during a scheduling window, enforced by (9), and (ii) the SINR of each cell i is higher than a given threshold, enforced by (10), in order to meet other QoS requirements discussed in section II.

$$\underset{\kappa,\mathbf{x}}{\text{maximize}} \quad c_\kappa(\mathbf{x}) \tag{8}$$

$$\text{subject to} \quad \sum_{k=1}^{\kappa} x_{i,k} = 1, \qquad \forall i \in \mathcal{C} \tag{9}$$

$$\sum_{k=1}^{\kappa} \sigma_i(\mathcal{C}_k(\mathbf{x}))\cdot x_{i,k} \geq B_i, \qquad \forall i \in \mathcal{C} \tag{10}$$

$$\kappa > 0,\ \mathbf{x} \in \{0,1\}^{C\times\kappa}$$

where $c_k(\cdot)$ is defined in (7), $\sigma_i(\cdot)$ is defined in (3), and $B_i \in \mathbb{R}_{\geq 0}$ is the minimum SINR threshold parameter.

This problem is NP-Hard. We prove NP-hardness by reduction from the graph coloring problem, which is known to be NP-hard [8].

In Problem 1, κ is a variable that influences the optimal solution through two coupled mechanisms : (i) it determines the fraction of the total capacity allocated to each time slot; and (ii) it affects the SINR of the cells assigned to the same time slot, i.e., $\{\mathcal{C}_k(\mathbf{x})\}_{k=1}^{\kappa}$. The NP-hardness of Problem 1, together with the nonlinear coupling between κ and the illumination plan $\mathbf{x}$, makes the computation of exact optimal solutions impractical for realistic network sizes.

A. A Graph Coloring Approximation

Problem 1 exhibits a close connection to the graph coloring problem. We now examine this connection more thoroughly and derive additional insights from it. Firstly, we introduce the *interference digraph* $\mathcal{D}(\mathbf{p}^{\text{Rx}},\alpha)$. It is parameterized by the received power matrix $\mathbf{p}^{\text{Rx}}$ and the sensitivity parameter $\alpha \in [0,1]$, such that

- the set of vertices is the same as the set of cells $\mathcal{C}$. For notational simplicity, vertex i refers to cell c_i in the rest of this section.
- the set of edges $\mathcal{A} \subseteq \mathcal{C}\times\mathcal{C}$ is constructed by selecting the top α fraction of most interfering cell pairs (i,j), i.e., those with the lowest pair-wise SIR:

$$(i,j)\in\mathcal{A} \iff \frac{p^{\text{Rx}}_{i,i}}{p^{\text{Rx}}_{i,j}} \leq Q_\alpha\left(\left\{\frac{p^{\text{Rx}}_{i,i}}{p^{\text{Rx}}_{i,j}}\right\}_{i\neq j}\right), \tag{11}$$

  where $Q_\alpha(\cdot)$ denotes the α-quantile of the empirical distribution of all off-diagonal pair-wise SIR values. When $\alpha=1$, all pairs are included and $\mathcal{D}(\mathbf{p}^{\text{Rx}},\alpha)$ is complete. When $\alpha=0$, no pairs are included and $\mathcal{D}(\mathbf{p}^{\text{Rx}},\alpha)$ is empty.

Therefore, $\mathcal{D}(\mathbf{p}^{\text{Rx}},\alpha)$ represents a mapping of the most significant pair-wise interference among cells. We can also represent the interference graph by the adjacency matrix $\mathbf{w}(\alpha)$ such that

$$w_{i,j}(\alpha) = \begin{cases} p^{\text{Rx}}_{i,j} & \text{if } (i,j)\in\mathcal{A} \\ 0 & \text{otherwise} \end{cases} \tag{12}$$

A heuristic way to approximate Problem 1 is to consider a graph coloring problem over the interference digraph. The idea is that each time slot can be seen as a different color, then we can re-use variables $x_{i,k}$ (hereby referred to as *coloring*) to assign cell i to color k. We consider C as an upper bound for the number of colors (the number of vertices). Moreover, we introduce variables

$$\mathbf{y} \in \{0,1\}^C \tag{13}$$

to indicate whether color k is used (i.e., $y_k=1$) or not (i.e., $y_k=0$). The resulting problem is formulated as follows

Problem 2. In this problem, our objective (14) is to find the chromatic number $\chi(\mathcal{D}(\mathbf{p}^{\text{Rx}},\alpha))$ for the interference digraph $\mathcal{D}(\mathbf{p}^{\text{Rx}},\alpha)$ subject to generalized conflict coloring constraints (17). We formalize it as follows:

$$\underset{\mathbf{x},\mathbf{y}}{\text{minimize}} \quad \sum_{k=1}^{C} y_k \tag{14}$$

$$\text{subject to} \quad \sum_{k=1}^{C} x_{i,k} = 1, \forall i\in\mathcal{C} \tag{15}$$

$$x_{i,k} \leq y_k, \forall i\in\mathcal{C}, k\leq C \tag{16}$$

$$\sum_{\substack{i\in\mathcal{C}\\ i\neq j}}\sum_{k=1}^{C} w_{i,j}(\alpha)\cdot x_{i,k}\cdot x_{j,k} \leq B_j, \forall j\in\mathcal{C} \tag{17}$$

$$\mathbf{x}\in\{0,1\}^{C\times\kappa},\ \mathbf{y}\in\{0,1\}^C$$

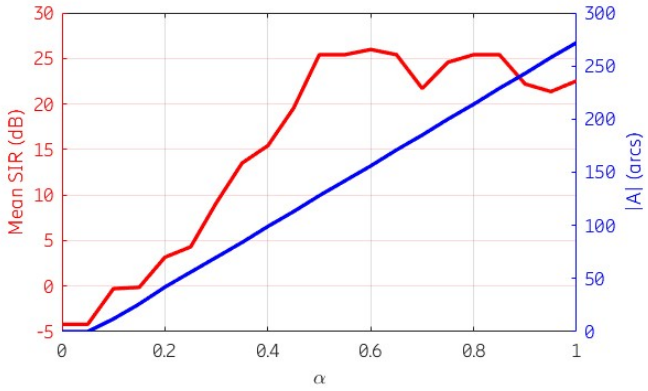


Fig. 3. Mean SIR and graph density as a function of α.

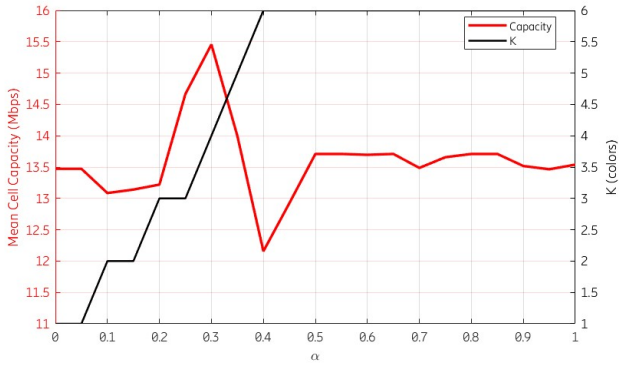


Fig. 4. Mean cell capacity and K as a function of α.

Similarly to Problem 1, (15) means that each cell is illuminated once and (16) ensures that a color is activated whenever it is assigned to at least one cell.

Remark 1. Problem 2's solution provides a good approximation for Problem 1 as the non-linear objective (8) can be implicitly maximized by minimizing the number of colors (14): Each color represents a time slot, and the total number of colors determines the scheduling cycle length. As the number of colors increases, each beam is served less frequently, and its time allocation decreases proportionally to $\frac{1}{K}$, where K is the number of used colors. The reduction in time allocation has a linear impact on the throughput, while the gain from improved SINR only grows logarithmically through the $\log(1 + \text{SINR})$ relationship. Consequently, the throughput degradation caused by increasing K cannot be fully compensated by SINR improvements, making color minimization a good approximation of the original objective.

B. Interference Graph's Properties

This subsection evaluates the impact of α on the scheduling performance. We consider a scenario with 17 cells arranged in a hexagonal pattern (more details in section IV). We sweep α from 0 to 1, with a step of 0.05. For each value, we construct $\mathcal{D}(\mathbf{p}^{\text{Rx}}, \alpha)$, apply the graph-coloring optimization to obtain the color assignment and number of colors K, then compute the per-cell SIR.

As shown in Figure 3, as α increases, the number of edges $|\mathcal{A}|$ increases linearly, meaning more cell pairs are identified as aggressors. These aggressors are taken into account by the scheduling optimization, which separates them into different time slots, thereby improving the SIR. However, beyond $\alpha = 0.5$, the SIR no longer improves: once the dominant interferers are already separated, the additional edges only capture weak interferers whose contribution to the total interference is negligible regardless of their scheduling.

Fig. 4 illustrates the mean cell capacity and the number of colors K as a function of α. For low values of α ($\alpha \in [0, 0.2]$), K remains low (from 1 to 3), meaning the optimization scheduler cannot fully separate the interferers with few available slots. As a result, some dominant aggressors remain co-scheduled, limiting the SIR and cell capacity improvement. For ($\alpha \in [0.2, 0.3]$), the cell capacity reaches its maximum of approximately 15.5 Mbps. Here, K is just large enough to separate the dominant interferers into different slots, and the SIR gain outweighs the bandwidth loss from dividing by K. At $\alpha \approx 0.4$, a sharp degradation occurs: the cell capacity drops to about 12.1 Mbps while K jumps to 6. At this point, the SIR has not yet improved enough to compensate for the increased number of colors, resulting in a cell capacity loss. For larger values ($\alpha \in [0.5, 1]$), the SIR improves and saturates as shown in Figure 3, but K remains high. The SIR gain partially recovers the capacity to around 13.5–13.7 Mbps, but the bandwidth division by K prevents it from reaching the optimal value observed at $\alpha \approx 0.3$. These results confirm the existence of an optimal operating point, in this scenario at $\alpha \approx 0.3$, and show that increasing α to include more aggressors does not necessarily improve performance. Beyond this point, adding more edges to the graph increases K without proportional SIR benefit, leading to fewer scheduled slots per cell.

IV. Experimental Setup

We run Matlab-based simulations where we consider a LEO satellite constellation at an altitude of 440 km with an orbital inclination of 53°, comprising 600 satellites distributed across 24 orbital planes (25 satellites per plane). Each satellite generates beams with a half-power beamwidth (HPBW) of 4.33°. The service link operates in S-band at a carrier frequency of 1 GHz with a channel bandwidth of 30 MHz. The deployment area covers 17 cells arranged in a hexagonal grid with an inter-site distance of 22.2 km. Each cell serves 0 to 50 randomly dispersed user terminals. The minimum elevation angle for the service link is set to 25°.

A. Baseline Algorithms

In our experiments, we compare the performance of three algorithms: (i) the solution provided by the classical *Branch-and-Bound* algorithm to solve Problem 2, which yields the minimum chromatic number required to achieve zero SIR constraint violations (17); (ii) a fixed regular Beam Hopping

Algorithm 1 Greedy Coloring Algorithm

Input: Graph $D = (\mathcal{C}, \mathcal{A})$, colors K
Initialize c randomly, $V \leftarrow \text{VIOLATIONS}(c)$
repeat
  $V_{\text{old}} \leftarrow V$
  for $i \in \mathcal{C}$ **do**
    $c_i \leftarrow \arg\min_{k \in \{1,\dots,K\}} \text{VIOLATIONS}(c_{i \leftarrow k})$
  end for
  $V \leftarrow \text{VIOLATIONS}(c)$
until $V = V_{\text{old}}$
return (c, V)

pattern that cyclically assigns x colors to cells based on their geographical positions, without interference-aware optimization. Hereafter, BH=x, with $x \in \{1, \dots, 4\}$, denotes this scheme with x colors (cells with the same colors are scheduled at the same time); and (iii) a *Greedy Heuristic* described in Algorithm 1, which, given a fixed number of colors, starts from a random coloring and iteratively improves it by testing all possible color reassignments for each cell and selecting the one that minimizes the total number of constraint violations (17), until no further improvement is found.

B. Performance Metrics

The following metrics are used to evaluate the beam scheduling algorithms:

- **Cell Capacity.** For a given coloring with K colors, the capacity of cell i is defined as:

$$c_i(\mathcal{C}') = \frac{\omega}{K} \cdot \log_2(1 + \sigma_i(\mathcal{C}')), \tag{18}$$

  where ω is the total channel bandwidth, K is the number of colors, the factor $1/K$ reflects the bandwidth reduction due to the resource sharing over time, and σ_i is computed as:

$$\sigma_i = \frac{1}{\frac{1}{\text{SIR}_i} + \frac{1}{\text{SNR}_i}} \tag{19}$$

- The SIR_i accounts only for co-slot interferers (i.e., cells assigned the same color as cell i)
- **Interference budget violations rate.** Number of nodes exceeding the interference budget.
- **User throughput satisfaction.** User throughput is evaluated using a birth–death queuing model, where arrivals correspond to UEs becoming active and departures to fully served UEs. For each cell, the arrival rate is $\lambda = \sum_{u=1}^{N_U} \lambda'_i(u)$, while the service rate $\mu(n)$ depends on the number of active users n.

V. Numerical results

This section presents the numerical results obtained from the NTN simulations. The experiments were executed using MATLAB's *intlinprog* solver.

Figure 5 presents the CDF of SIR at cell level, aggregated over all time steps. As expected, the baseline scenario BH=1, where all beams are simultaneously active, yields the lowest SIR due to maximum inter-beam interference. For regular beam hopping schemes, increasing the number of colors (i.e., time slots) reduces the number of co-channel beams per slot, thereby improving SIR. Similarly, for the optimization-based approach, increasing α leads to a denser aggressor graph, resulting in a higher chromatic number K and consequently higher SIR values.

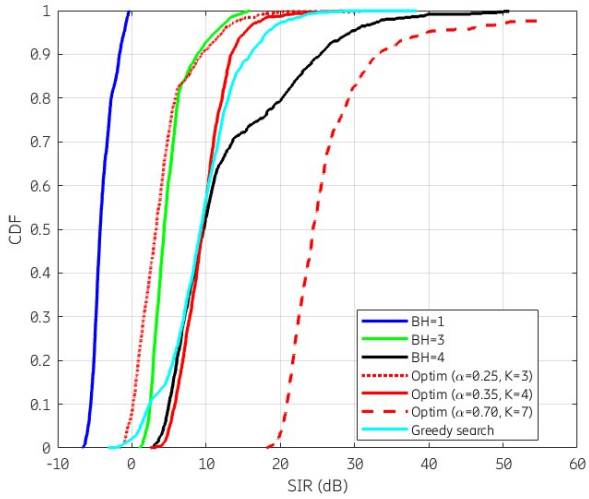


Fig. 5. SIR cdf comparison for all techniques.

When comparing the optimization approach with regular beam hopping at the same number of colors, both methods achieve comparable average SIR performance. However, at $K = 4$, a significant fairness gap emerges between both approaches. We can observe this in the SIR distributions, where the optimization with $\alpha = 0.35$ exhibits a steeper SIR CDF, while the BH=4 shows a broader spread of SIR values. This indicates that regular beam hopping, which only relies on geometrical patterns, shows high SIR improvement for best-served users that are generally located at the cell centers whereas the graph-coloring optimization produces a more balanced allocation for all users, even the ones at the edge. Finally, the greedy heuristic algorithm with 4 colors achieves performance comparable to the ILP-based optimization with 4 colors, offering a computationally efficient alternative for practical deployment scenarios.

Figure 6 presents the interference budget constraint violations for each scheduling strategy. The baseline BH=1 exhibits 100% violations, as all beams are simultaneously active. Regular beam hopping BH=3 reduces violations to 70.6%, but remains far from satisfying the constraints. All three optimization variants achieve 0% violations by construction, regardless of the value of α. The greedy heuristic achieves 5.9% violations, representing a significant improvement over regular beam hopping with 3 colors while maintaining low computational complexity.

Table I details the per-UE capacity statistics : mean, median, percentiles 5% and 95% and Jain Fairness Index.

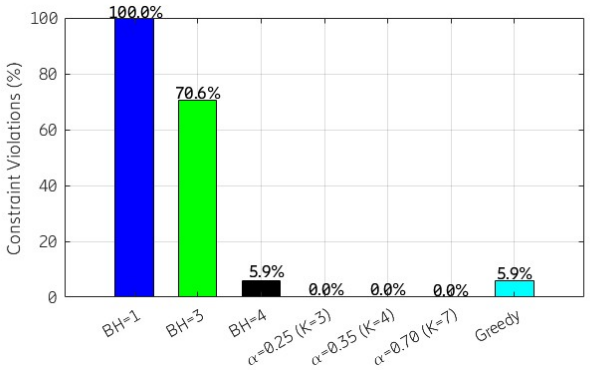


Fig. 6. Interference Budget constraint violation.

TABLE I
USER CAPACITY (MBPS)

Algorithm	Mean	Median	P5	P95	JFI
BH=1	0.48	0.41	0.04	1.03	0.5772
BH=3	0.51	0.42	0.04	1.25	0.5443
BH=4	0.52	0.43	0.04	1.31	0.5473
Optim ($\alpha = 0.35$, $K = 4$)	0.52	0.44	0.04	1.18	0.5623
Greedy search	0.51	0.43	0.04	1.19	0.5545

As shown in Table I, the optimization with $\alpha = 0.35$ achieves the highest mean capacity (0.52 Mbps), the highest median (0.44 Mbps), and the best Jain's Fairness Index (0.5623) among all evaluated schemes. Compared to the regular BH=4, which achieves the same mean but a higher P95, the optimization provides better fairness and median performance. While BH=4 benefits the best-served users, it does so at the expense of overall fairness. The greedy heuristic performs comparably to BH=3.

Table II reports the user satisfaction (served/offered) for the considered configurations. The optimized scheme (opti $K = 4$) achieves the highest mean and median values for this scenario, indicating an improvement in the typical user experience compared to the baseline BH=4 configuration.

From an implementation perspective, the proposed optimization relies on an ILP formulation derived from a graph coloring problem, and is therefore NP-hard. When solved using MATLAB's *intlinprog* solver based on a branch-and-bound framework, the worst-case computational complexity is exponential in the number of integer variables, which grow with both the number of cells and colors. While solving large instances to optimality may become intractable, the computational burden can be controlled through solver parameters such as *MaxTime*, *MaxNodes*, and *RelativeGapTolerance*. Furthermore, since satellite trajectories are predictable, the optimization can be performed offline at the Satellite Operator Center (SOC), eliminating the need for real-time onboard execution.

VI. CONCLUSION

This paper proposed an interference-aware beam hopping framework for LEO satellite networks based on soft-conflict graph coloring. By modeling interference through a parameterized conflict graph, the approach jointly manages interference and temporal resource allocation. Unlike conventional beam hopping schemes based on fixed geometric patterns, it exploits physical interference information, making it adaptable to heterogeneous NTN deployments. Numerical results show comparable performance to regular beam hopping schemes, with slight user throughput satisfaction gains in some scenarios. They also highlight a range of values for the interference selection parameter α that balances interference mitigation and time-slot fragmentation. More broadly, the framework provides a generic foundation for NTN resource management and can be extended to support irregular cell layouts, heterogeneous traffic demands, multi-satellite coordination, and adaptive interference management. Future work will investigate dynamic graph construction, learning-based tuning of α, and traffic-aware scheduling.

TABLE II
USER TRAFFIC SATISFACTION (SERVED/OFFERED)

Algorithm	Mean	Median
BH=4	0.452	0.485
Optim K=4	0.460	0.530
BH=3	0.443	0.455

REFERENCES

[1] J. Anzalchi et al., "Beam hopping in multi-beam broadband satellite systems: System simulation and performance comparison with non-hopped systems," *2010 5th ASMS Conference and the 11th Signal Processing for Space Communications Workshop*, 2010.

[2] Y. Wang et al., "A flexible resource allocation algorithm in full bandwidth beam hopping satellite systems," *2019 IEEE 3rd Advanced Information Management, Communicates, Electronic and Automation Control Conference (IMCEC)*, 2019.

[3] R. Mu et al., "Dynamic beam hopping for leo satellites with differentiated traffic demands," *WSCP2023*, 2023.

[4] H. Deng et al., "Satellites beam hopping scheduling for interference avoidance," *IEEE Journal on Selected Areas in Communications*, 2024.

[5] Z. Liu et al., "Multi-satellite coordinated beam hopping for interference mitigation under tilted beam effects: A graph-theoretic approach," *IEEE Wireless Communications Letters*, 2026.

[6] H. T. Friis, "A note on a simple transmission formula," *Proceedings of the IRE*, 1946.

[7] 3rd Generation Partnership Project (3GPP), "Study on new radio (nr) to support non-terrestrial networks (release 15)," 3GPP, 2018.

[8] M. R. Garey et al., *Computers and Intractability: A Guide to the Theory of NP-Completeness*. New York: W. H. Freeman and Company, 1979.