Robust Beam-to-Satellite Routing Strategies for Megaconstellations

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

***Abstract*—Robust routing strategies for satellite constellations are crucial for mitigating outage probability amidst satellite failures. While existing literature primarily addresses robust routing within the satellite network, this study introduces innovative strategies to enhance robustness in ground-satellite links. Leveraging established heuristic and optimization methods, these strategies are tailored to incorporate redundancy. Using the Starlink constellation with 200,000 users as an example, our findings illustrate a significant reduction in outage probability for end-users, albeit with a trade-off of reduced capacity. Specific approaches demonstrate the capability to sustain a substantial portion of throughput, maintaining 96% while reducing the outage probability up to 66% and 10% for internal and external failures, respectively.**

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

I. INTRODUCTION

Satellite communication constellations are currently witnessing an unprecedented surge in capacity growth [1], presenting a viable broadband connectivity solution for regions with limited terrestrial network coverage. This expansion is facilitated by heightened spacecraft flexibility, made possible by technological advancements, enabling dynamic adaptation to varying user demands. Leveraging this newfound flexibility necessitates the development of advanced resource allocation algorithms responsible for distributing constellation resources, including spectrum and power, marking an incipient area of research [2], [3]. Among these resources, optimized routing strategies play a pivotal role in enhancing satellite utilization by efficiently distributing workload across satellites [4].

In recent literature, a variety of methodologies have emerged to address the satellite routing problem, which entails mapping users to satellites across various conditions. Within broadband networks, existing approaches leverage a range of techniques including greedy algorithms [5], deep reinforcement learning (DRL) [6], sub-gradient optimization [7], and mixed-integer linear programming [4] to effectively distribute workload among satellites, thereby enhancing capacity in high-dimensional environments with potential interference. Moreover, the satellite routing problem has garnered significant attention beyond broadband communication, extending into contexts such as mobile networks [8] and task-oriented networks [9] (e.g., imaging satellites).

Furthermore, there has been a growing interest in exploring the robustness and reliability of satellite constellation routing. Reference [10] delves into mechanisms aimed at mitigating satellite overloading, which can lead to temporary satellite failures in densely populated regions. Reference [11] introduces a DRL-based approach integrating jamming mitigation strategies to enhance overall performance. Reference [12] showcases the resilience of existing satellite constellation designs against geographical failures in inter-satellite links. Nevertheless, existing research mainly focuses on enhancing robustness within the satellite network itself, largely overlooking ground-satellite connectivity. Given the potential for satellite failures arising from internal or external factors such as component malfunction, solar storms, or jamming, there arises a pressing need to develop robust satellite routing algorithms that incorporate redundancy within ground-satellite links.

This paper aims to bridge this gap by: 1) introducing robust satellite routing strategies that incorporate redundancy, and 2) evaluating their performance across various failure scenarios. This work builds upon the methods in [4], which have been adapted to integrate redundancy. Our results showcase that these robust strategies markedly diminish outage probability in the event of satellite failure, albeit with a reduction in capacity. Specific strategies are capable of sustaining 96% of throughput, while reducing the outage probability by up to 66% and 10% for internal and external failures, respectively.

II. THE SATELLITE ROUTING PROBLEM

The satellite routing problem involves determining the optimal mapping between beams and satellites at each time instance. To achieve efficient solutions, it is imperative to balance the load across satellites while minimizing interference [4]. Formally, a set of beams $\mathcal{B}$, each with demand d_b, needs to be mapped to a corresponding set of satellites $\mathcal{S}$. At any given time, each beam has a subset of visible satellites $\mathcal{S}_b^*$ ordered by priority. To achieve a solution, each beam must be assigned to one satellite within its subset of visible satellites. We introduce a binary variable $x_{b,s}$ to signify whether beam b is mapped to satellite s at a specific time. By enforcing the constraint: $C1 : \sum_{s\in\mathcal{S}_b^*} x_{b,s} = 1 \ \forall b \in \mathcal{B}$, we ensure that each beam is assigned to exactly one satellite. A solution to the satellite routing problem is valid if constraint $C1$ is satisfied.

To tackle this challenge, we leverage the methodologies outlined in [4], briefly summarized here. Initially, the authors introduce a maximum elevation angle heuristic as a baseline solution, denoted as ME-1, wherein each beam b is linked to the satellite with the highest elevation angle.

Manuscript received MONTH DAY, 2024; 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 FORMULATIONS FOR THE SATELLITE ROUTING RESOLUTION ACCORDING TO [4]

Phase	Formulation	General and Variable description
Beam clustering (solved once)	$\min_{x_{b,c}} \quad \sum_{b_1 \in \mathcal{B}, b_2 \in \mathcal{B}} (f_{b_1,b_2} y^A_{b_1,b_2})$ $s.t. \quad C2: y^A_{b_1,b_2} \geq z_{b_1,b_2} \mathbb{1}_{\{c_1,c_2\} \in \mathcal{R}_B} (x_{b_1,c_1} + x_{b_2,c_2} - 1)$ $C3: \sum_{c \in \mathcal{C}_b} x_{b,c} = 1$ (1)	**Objective:** Assign nearby beams to different clusters to increase balance across satellites **Variables**: $x_{b,c}$: Mapping of beam b to cluster c $y^A_{b_1,b_2}$: Indicates if b_1 and b_2 share a satellite at some point in time z_{b_1,b_2}: Indicates if b_1 and b_2 observe the same satellite at some point in time $\mathcal{R}_B$: Set of equivalent clusters **Constraints**: $C2$: Activate $y^A_{b_1,b_2}$ on overlapping satellites $C3$: Ensure each beam is mapped to a cluster
Cluster to satellite (solved for each time t)	$\min_{x_{s,c}} \quad \sum_{b_1 \in \mathcal{B}, b_2 \in \mathcal{B}} y^E_{b_1,b_2}$ $s.t. \quad C4: y^E_{b_1,b_2} \geq I_{b_1,s_1,b_2,s_2} (x_{b_1,s_1} + x_{b_2,s_2} - 1)$ $C5: x_{b,s_i} = \sum_{\{c_b,c\} \in \mathcal{R}_D} (x_{s_i,c} - \sum_{j=1}^{i-1} x_{s_j,c})$ (2) $C6: \sum_{\{c_b,c\} \in \mathcal{R}_D} \sum_{s \in \mathcal{S}^*_b} x_{s,c} \geq 1$ $C7: \sum_{c \in \mathcal{C}} x_{s,c} = 1$	**Objective**: Assign clusters to satellites while minimizing interference **Variables**: $x_{b,s}$: Mapping of beam b to satellite s $x_{s,c}$: Mapping of satellite s to cluster c $y^E_{b_1,b_2}$: Indicates if b_1 and b_2 have potential interference at time t I_{b_1,s_1,b_2,s_2}: Indicates potential interference between b_1 served by s_1 with b_2 served by s_2 at time t $\mathcal{R}_D$: Set of compatible clusters **Constraints**: $C4$: Activate $y^E_{b_1,b_2}$ on potential interference $C5$: Map beam b to satellite s, if feasible $C6$: Ensure each beam has a visible satellite with matching cluster $C7$: Map each satellite to a cluster

Nevertheless, the ME-1 method demonstrates notable limitations, particularly in densely populated regions where beams are geographically close and in high demand. To address this challenge, the authors propose a two-phase approach based on integer programming, denoted as IP-1. In the initial IP-1 phase, termed *beam clustering*, beams are aggregated into clusters, aiming to disperse nearby beams across different clusters. Subsequently, in the second IP-1 phase, termed *cluster to satellite*, these clusters are allocated to satellites with the objective of minimizing interference both within and between satellites. The formulation and description of each phase are outlined in Table I. Following resolution, both the ME-1 and IP-1 methods are used to identify pairs of beams sharing a satellite at any given time (grouped in set $\mathcal{R}_A$) and pairs potentially causing interference (grouped in set $\mathcal{R}_E$). These sets are then fed into the Frequency Assignment problem to ensure appropriate spectrum allocation. By dispersing nearby beams across different satellites while mitigating interference, the IP-1 method yields a capacity increase of up to 75% in constellations like Starlink compared to the ME-1 method.

III. ROBUST SATELLITE ROUTING STRATEGIES

This study introduces three robust strategies that build upon the framework presented in [4]: one strategy based on the ME heuristic and two strategies based on the IP technique. These strategies are designed to yield solutions with more conservative $\mathcal{R}_A$ and $\mathcal{R}_E$ sets, thereby resulting a constrained frequency assignment and freeing up resources on the satellites. During operations, these redundancies empower users to select alternative satellites with available resources. The three strategies, alongside the original methods, are illustrated in Figure 1. It is noteworthy that these strategies complement other robust resource allocation approaches, such as the Frequency Assignment strategies outlined in [13].

A. Redundant heuristic: ME-X

The robust heuristic strategy, denoted as ME-X, entails mapping each beam to multiple satellites rather than just one during the computation of the $\mathcal{R}_A$ and $\mathcal{R}_E$ sets. This approach expands the scope of beams potentially sharing a satellite or experiencing interference, leading to a more conservative Frequency Assignment. Specifically, instead of assigning each beam to the satellite with the maximum elevation angle, we assign each beam to the X satellites with the highest elevation angles. During operations, each beam is linked to the satellite with the highest elevation angle that remains operational. If all X satellites with the highest elevation angle have failed, it is assumed that the beam cannot be serviced. Note that, when X=1, the strategy is equivalent to the baseline heuristic ME-1.

B. Redundant optimized: IP-X

The first robust optimized strategy, denoted as IP-X, adopts a similar concept to the robust heuristic, introducing redundancy by providing a set of X compatible satellites instead of only one. In the formulation, we ensure that each beam has at least X_b satellites with a compatible cluster. This constraint can be expressed as:

$$C6': \sum_{\{c_b,c\} \in \mathcal{R}_D} \sum_{s \in \mathcal{S}^*_b} x_{s,c} \geq X_b \qquad (3)$$

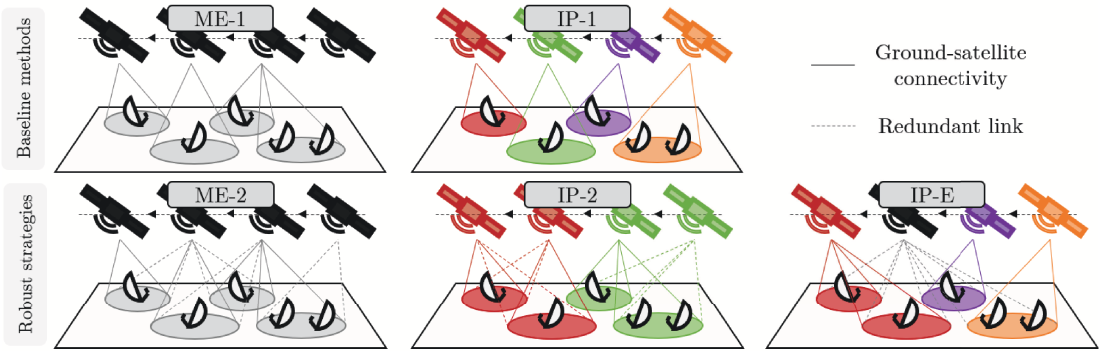

Fig. 1. Depiction of robust strategies formulated in this work

Here, X_b is constrained to be the minimum between X and the number of visible satellites for beam b. Subsequently, we enforce that each beam is mapped to at least X_b satellites. This is achieved by introducing an additional binary variable n_{b,s_i}, equal to 0 when the number of satellites with a compatible cluster with higher priority (i.e., appearing before satellite s_i in set $\mathcal{S}_b^*$) is less than X, 1 otherwise. If n_{b,s_i} equals 0 and s_i has a compatible cluster, beam b must be mapped to satellite s_i. The new constraints are as follows:

$$\begin{aligned} C5' : x_{b,s_i} &= \sum_{\{c_b,c\}\in\mathcal{R}_D} x_{s_i,c} - n_{b,s_i} \\ C8 : -\sum_{\{c_b,c\}\in\mathcal{R}_D} \sum_{j=1}^{i-1} x_{s_j,c} &+ X n_{b,s_i} \leq 0 \end{aligned} \tag{4}$$

By incorporating $C8$ and substituting $C5$ with $C5'$ and $C6$ with $C6'$, the formulation is adapted to integrate redundancy. Solving this formulation using mathematical solvers and the iterative process outlined in [4] yields $\mathcal{R}_A$ and $\mathcal{R}_E$ sets that consider redundancy at the beam level. It is important to note that during operations, the original formulation remains unchanged. Furthermore, in line with the approach described in [4], if two beams are potentially interfering during operations, one of them is deactivated as a precautionary measure. Note that, when X=1, the strategy is equivalent to the baseline optimization IP-1.

C. Emergency satellite: IP-E

The objective of the second robust strategy, denoted as IP-E, is to introduce redundancy at the cluster level, aiming to enhance utilization while ensuring robustness. To achieve this, we introduce a new cluster, denoted as c_E, representing an emergency satellite. Each beam must have at least one visible satellite assigned to the c_E cluster. In the event of a satellite failure, all beams mapped to the original satellite are redirected to the emergency one. This constraint is expressed as follows:

$$C9 : \sum_{s\in\mathcal{S}_b^*} x_{s,c_E} \geq 1 \tag{5}$$

Note that this constraint applies only to beams with more than one visible satellite. By incorporating $C9$ into the original formulation, redundancy at the cluster level is integrated. During operations, the original formulation remains unchanged.

IV. RESULTS

The experimental setup of this study closely resembles that described in [4]. The constellation configuration and user antenna characteristics align with those of the Starlink system, with values sourced from public filings. User distribution comprises 20,000 locations computed based on population distribution [14], with each location assumed to host 10 users, each with a consistent demand of 100 Mbps. All other parameters are extracted from [4].

To evaluate robustness against satellite failure, we consider five scenarios where 0, 20, 50, 200, and 500 satellites respectively fail simultaneously during operations. Two types of failures are examined: **random** failures, where the selected failing satellites are distributed randomly, and **localized** failures, where one satellite is randomly chosen, and the nearest satellites, including the selected one, are assumed to fail. The former simulates internal failures (e.g., component issues), while the latter represents external failures (e.g., solar storms). Each scenario is simulated for 50 time-steps, with a 120-second interval between time-steps. To enhance confidence, each simulation is repeated 10 times with different sets of failing satellites. The commercial solver Gurobi [15] (version 9.1.2) is employed for optimization tasks on an Intel(R) Xeon(R) Platinum 8160 CPU @ 2.10GHz, with up to 16 simultaneous threads permitted.

The results concerning average throughput and outage probability over time, computed as the percentage of unserved users, are depicted in Figure 2. Several observations can be made. First, the baseline heuristic method (ME-1) exhibits a linear decrease in throughput and rise in outage probability proportional to the number of failed satellites, with the reduction in throughput and increase in outage reaching around 20% when 1000 satellites failed. Including robustness into this strategy (ME-2) mitigates outage probability to 6% and 10% on 1000 random and localized satellite failures, respectively. Augmenting robustness further (ME-3) yields additional reductions in outage probability to 1% and 7%, respectively. Note that including robustness into the heuristic strategy does not compromise throughput, leading to an overall superior performance.

Turning to the baseline optimization method (IP-1), we observe that, even when not adapted for robustness, the outage probability remains comparable to that achieved with ME-2.

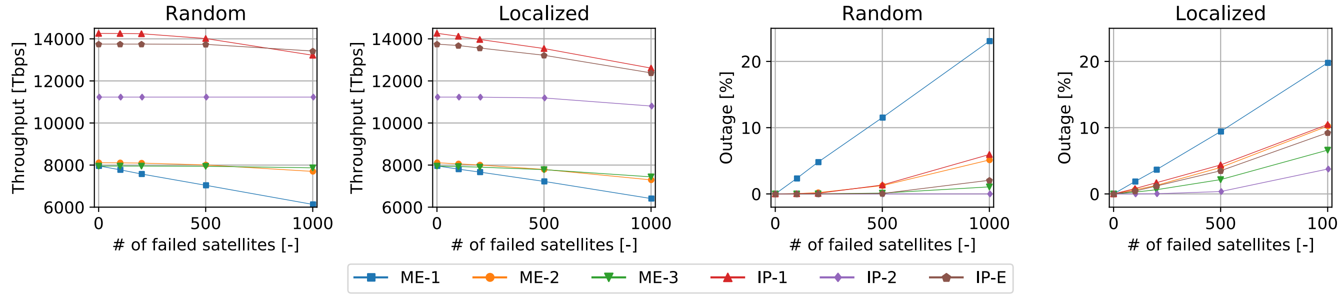

Fig. 2. Results on the six strategies on the SpaceX constellation on random and localized satellite failures

This suggests that the inherent optimization process possesses some capacity to address satellite failures to a certain extent. Including robustness at the beam level (IP-2) reduces outage probability to 0% and 4% for the random and localized scenarios, respectively. However, this strategy comes at the expense of throughput, with a 21% reduction even when all satellites operate in nominal conditions, highlighting the trade-off inherent in pursuing robustness. Adding robustness at the cluster-level (IP-E) yields a substantial reduction in outage probability to 2% and 9% for the random and localized scenarios, respectively, while having a reduced impact on throughput, being only 4% lower than the original method. Note that this translates to a reduction of up to 66% in outage probability compared to the original IP-1 method on random failures, and up to 10% on localized failures. Finally, it is observed that random failures exert a lesser impact on outage probability compared to localized failures, with the exception of the ME-1 strategy.

As a final remark, note that simultaneous failures, while rare, do happen occasionally. As example, a solar storm in 2022 destroyed 38 satellites [16], while a design flaw forced SpaceX to decommission 100 satellites earlier this year [17]. It is expected that, as the number of satellites goes up, these events will increase in frequency, making robust strategies crucial to minimize outage probability.

V. CONCLUSIONS

In this study, we have introduced and examined robust strategies to address the Satellite Routing problem, focusing on the allocation of beams to satellites. These strategies, which extend existing literature, introduce redundancy on the satellite-user links. When simulating the proposed methods using the SpaceX constellation with 200,000 users, our findings indicate a significant decrease in outage probability across both random and localized failure scenarios. When using optimized strategies, introducing redundancy comes at the cost of reduced throughput. Nonetheless, certain strategies demonstrate the capacity to mitigate outage probability by up to 66% and 10% for internal and external failures, respectively, with a minimal throughput reduction of merely 4% in comparison to methods without redundancy.

ACKNOWLEDGMENTS

The authors would like to thank SES for their financial support.

REFERENCES

[1] 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.

[2] C. N. Efrem and A. D. Panagopoulos, “Dynamic energy-efficient power allocation in multibeam satellite systems,” *IEEE Wireless Communications Letters*, vol. 9, no. 2, pp. 228–231, 2019.

[3] M. Takahashi, Y. Kawamoto, N. Kato, A. Miura, and M. Toyoshima, “Adaptive power resource allocation with multi-beam directivity control in high-throughput satellite communication systems,” *IEEE Wireless Communications Letters*, vol. 8, no. 4, pp. 1248–1251, 2019.

[4] N. Pachler, E. F. Crawley, and B. G. Cameron, “Avoiding Self-Interference in Megaconstellations through Cooperative Satellite Routing and Frequency Assignment,” *IEEE Journal on Selected Areas in Communications*, 2024.

[5] J. Li, K. Xue, J. Liu, and Y. Zhang, “A user-centric handover scheme for ultra-dense leo satellite networks,” *IEEE Wireless Communications Letters*, vol. 9, no. 11, pp. 1904–1908, 2020.

[6] 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.

[7] 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.

[8] 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.

[9] F. Xhafa and A. W. Ip, “Optimisation problems and resolution methods in satellite scheduling and space-craft operation: a survey,” *Enterprise Information Systems*, vol. 15, no. 8, pp. 1022–1045, 2021.

[10] L. Zhang and Y. Du, “A dynamic cascading failure model for leo satellite networks,” *IEEE Transactions on Network and Service Management*, 2023.

[11] K. Chu, S. Cheng, and L. Zhu, “A robust routing strategy based on deep reinforcement learning for mega satellite constellations,” *Electronics Letters*, vol. 59, no. 11, p. e12820, 2023.

[12] Q. Ouyang, N. Ye, S. Miao, B. Kang, A. Wang, and L. Zhao, “Mega constellation networks are reliable against geographical failure,” in *2023 IEEE 98th Vehicular Technology Conference (VTC2023-Fall)*. IEEE, 2023, pp. 1–5.

[13] G. Casadesus, J. J. G. Luis, N. Pachler, E. F. Crawley, and B. G. Cameron, “Dynamic frequency assignment for mobile users in multibeam satellite constellations,” in *International Astronautical Congress*, 2022.

[14] 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.

[15] Gurobi Optimization, LLC, “Gurobi Optimizer Reference Manual,” 2023. [Online]. Available: https://www.gurobi.com

[16] T. Dang, X. Li, B. Luo, R. Li, B. Zhang, K. Pham, D. Ren, X. Chen, J. Lei, and Y. Wang, “Unveiling the space weather during the starlink satellites destruction event on 4 february 2022,” *Space weather*, vol. 20, no. 8, p. e2022SW003152, 2022.

[17] SpaceX, “Commitment to space sustainability,” https://api.starlink.com/public-files/Commitment%20to%20Space%20Sustainability.pdf, 2024, accessed: 2024-05-01.