

Assessing LEO Satellite Networks for National Emergency Failover

Vaibhav Bhosale
Georgia Institute of Technology
Atlanta, USA

Robin Kim
Georgia Institute of Technology
Atlanta, USA

Ekaterina Tumanova
Georgia Institute of Technology
Atlanta, USA

Alberto Dainotti
Georgia Institute of Technology
Atlanta, USA

Ying Zhang
Northwestern University
Evanston, USA

Miguel Schlicht
Georgia Institute of Technology
Atlanta, USA

Zachary S. Bischof
Georgia Institute of Technology
Atlanta, USA

Ahmed Saeed
Georgia Institute of Technology
Atlanta, USA

Sameer Kapoor
Georgia Institute of Technology
Atlanta, USA

Muskaan Gupta
Georgia Institute of Technology
Atlanta, USA

Fabián E. Bustamante
Northwestern University
Evanston, USA

Abstract

In this paper, we study the viability of LEO networks as a failover network. We contextualize our analysis by framing the capacity of satellite networks relative to lost capacity due to submarine cable failure. Specifically, we focus on scenarios where LEO networks act as failovers for submarine cables, providing a concrete target capacity to be fulfilled by the satellite network. We introduce a new model and simulator that help us estimate the failover capacity. We identify key factors determining the actual capacity available on the satellite network: the total area of the country, the terminal distribution policy used by the government, the spectrum allocation and traffic engineering policies used by the LEO network operator. Based on our findings, we make policy recommendations to governments that can result in an increase of up to $1.8\times$ in the failover capacity without requiring additional infrastructure. However, we find after implementing all our recommendations, with 200k terminals deployed and no competing traffic in the network, a satellite network can only satisfy 0.9-14.7% of the capacity lost due to submarine cable failure in four out of six case studies.

CCS Concepts

• **Networks** → **Network simulations**; **Network performance modeling**; *Network measurement*.

Keywords

Low Earth Orbit Satellite Networks; Network Capacity Estimation

ACM Reference Format:

Vaibhav Bhosale, Ying Zhang, Sameer Kapoor, Robin Kim, Miguel Schlicht, Muskaan Gupta, Ekaterina Tumanova, Zachary S. Bischof, Fabián E. Bustamante, Alberto Dainotti, and Ahmed Saeed. 2025. Assessing LEO Satellite Networks for National Emergency Failover. In *Proceedings of the 2025 ACM Internet Measurement Conference (IMC '25)*, October 28–31, 2025, Madison, WI, USA. ACM, New York, NY, USA, 20 pages. <https://doi.org/10.1145/3730567.3764482>

1 Introduction

Low Earth Orbit (LEO) satellite networks are rapidly expanding, offering global coverage and playing an increasingly important role in Internet resilience. Several providers have deployed thousands of satellites, with plans for many more. LEO systems have already proven reliable during infrastructure outages [21, 58, 74, 79]. Their potential for enhancing resilience has prompted governments to adopt more systematic strategies by a number of different countries including Ghana [36], Israel [76], Taiwan [19] as well as efforts like NATO's satellite-based backup for submarine cable failures [7].

The performance of a LEO network in emergency scenarios depends heavily on the agreement between the government and the network operator, raising the question: how should such agreements be structured? Answering this requires navigating a multifaceted landscape, with constraints across the network stack (e.g., spectrum allocation, traffic engineering) and multiple stakeholders with differing priorities. These include local governments, existing LEO users, and operators themselves, each weighing deployment strategy, political sensitivities (e.g., sovereignty and data access), and service quality for current customers.

LEO network viability cannot be assessed in isolation. Its value varies widely depending on perspective. For instance, deploying a few Starlink terminals to support first responders during a hurricane can be transformative. But whether that benefit scales to support an entire city or a nation remains unclear. LEO networks today do not yet demonstrate elastic capacity at such scales, and large-scale deployment often requires coordination far beyond local initiatives.



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

IMC '25, October 28–31, 2025, Madison, WI, USA

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1860-1/2025/10

<https://doi.org/10.1145/3730567.3764482>

		Tonga	Haiti	Lithuania	Ghana	Great Britain	South Africa
§4	Lost Capacity because of Cable Failure	320 Gbps	320 Gbps	101 Gbps	83,700 Gbps	400,000 Gbps	31,700 Gbps
§5.2	LEO Network Max Capacity (% of cable capacity)	41 Gbps (12.8%)	1,389 Gbps (434%)	2,005 Gbps (198%)	2,163 Gbps (2.6%)	3,530 Gbps (0.9%)	4,653 (14.7%)
	Impact of Population Density on Network Capacity	Moderate (50%)	Significant (87.5%)	Significant (81%)	Significant (83%)	Moderate (51%)	Significant (78%)
	# Terminals (for 90% of Max Capacity)	500	20,000	20,000	50,000	50,000	100,000
§5.2 §5.4	Bottlenecks	RF	RF	RF	Satellite Count	Satellite Count	Satellite Count
§5.5	Impact of Failover Traffic on Global Coverage	Minimal	Minimal	Local	Local	Far Reaching	Far Reaching

Table 1: Highlighting the maximum capacity a LEO network (Starlink) can provide (in comparison to lost submarine cable capacity) along with number of terminals required to achieve 90% of that capacity, bottlenecks to LEO network capacity, impact of population sparsity, and the impact on global coverage for the six case studies. Population sparsity impact refers to the impact of population being distributed in smaller pockets of the nation’s landmass. RF bottleneck can be alleviated through more efficient spatial multiplexing (i.e., narrower beams and smaller coverage cells) or the allocation of more bandwidth.

Compounding this is the opaqueness of LEO system design and operations, which continue to evolve rapidly as new technology and regulatory frameworks emerge [5]. Governments need better tools and data to inform infrastructure decisions. Likewise, researchers need models that bridge existing knowledge in capacity planning for terrestrial networks (e.g., cellular and submarine systems) with the novel properties of LEO networks.

This paper takes a first step by evaluating the viability of LEO networks as national-scale failover for submarine cable disruptions. National governments, given their reach and bargaining power, are well-positioned to define failover requirements and negotiate network behavior [14, 19, 29, 36, 52], offering a natural setting to explore cooperative and systematic resilience planning.

Despite the availability of LEO network simulators [26, 44, 45, 50], none support scalable analysis of aggregate network capacity under varying terminal distributions, Radio Frequency (RF) resource allocation, and traffic engineering policies. Most focus on individual link performance and do not scale to tens of thousands of satellites and hundreds of thousands of terminals. They also often overlook RF allocation constraints critical to understanding real-world capacity limits. To address this gap, we introduce CosmoSim, a lightweight, programmable, and modular simulator for studying the aggregate capacity of LEO networks across diverse deployment and policy scenarios.

Our approach combines empirical data, a realistic simulation model, and real-world submarine cable failures. Our analysis focuses on Starlink, the largest deployed LEO network, as it provides the only example of a network operating at a scale relevant to our study. Further, our capacity estimates represent an optimistic upper bound of the failover capacity provided by Starlink, ignoring many possible inefficiencies in operations and deployment (e.g., poor weather condition and unavailability of line-of-sight between user terminals and satellites). However, our simulator can be extended to capture the behavior of any LEO constellation and incorporate more realistic network conditions. We focus on six case studies, countries recently affected by cable outages, selected based on factors that influence satellite capacity, such as land area and gateway

presence. To contextualize our findings and motivate policy implications, we compare the capacity lost during each outage with the capacity Starlink could offer under different deployment scenarios.

Our model identifies four key factors that determine the capacity available in a satellite network: the country’s total area and population distribution; the deployment of user terminals; the operator’s spectrum allocation policy; and its traffic engineering strategy. We also evaluate how capacity may evolve as LEO constellations expand and wireless link performance improves.

Our analysis highlights critical trade-offs that governments and network operators must consider when planning for emergency connectivity. These include the effects of sovereignty-based restrictions, terminal deployment strategies, and national spectrum allocation policies. These insights offer a foundation for cooperation frameworks between governments and satellite providers. A summary of our findings appears in Table 1.

Our key contributions are the following. *First*, we frame the problem of evaluating the aggregate capacity of satellite networks as national-scale failover infrastructure in the context of submarine cable failures. To support this, we develop a new model and simulator, which is available at <https://github.com/GT-ANSR-Lab/CosmoSim>. *Second*, we define the terminal distribution problem as a cornerstone of maximizing the capacity of failover satellite networks, and present an efficient heuristic for solving it. *Third*, our analysis yields ten key insights into factors that limit LEO network capacity in failover scenarios. For example, distributing terminals solely based on population density can lead to severe RF contention, reducing available capacity by up to 50% in the case of South Africa. *Finally*, we provide three policy recommendations for governments managing their spectrum resources and cooperating with satellite network operators. For example, jointly optimizing terminal placement and RF allocation can improve available capacity by 1.8×.

2 Background

As LEO satellite networks become increasingly integrated into global communications, their role in emergency connectivity and national resilience planning has gained significant attention. Beyond their use in remote areas, LEO networks have demonstrated

their ability to provide critical connectivity during infrastructure failures, supporting disaster response, military operations, and large-scale outages. Given this emerging role, national governments and international organizations are exploring ways to systematically incorporate LEO networks as a failover mechanism for terrestrial infrastructure. However, deploying and regulating these networks presents unique challenges, including licensing, spectrum allocation, and capacity constraints imposed by shared satellite infrastructure. This section examines the policies surrounding LEO network deployment, focusing on their use as a failover solution and the regulatory frameworks that govern their operation. We close the section with a brief overview of related work.

2.1 LEO Networks Policies

LEO networks as failover. Starlink already has a high-profile track record of supporting critical operations when all other network infrastructure fails, including supporting first responders in Florida [74] and North Carolina [79], protesters in Iran [63, 88], and military and civilian operations in Ukraine [21, 58], where over forty thousand terminals are deployed to provide secure and reliable communication [89]. In all these scenarios, the deployment of satellite networks was reactive. There are also proactive proposals for dealing with outages using satellite networks [38, 43]. For example, Starlink recently started offering a low-cost “backup plan” where inactive users pay a low monthly fee to use Starlink when the user’s primary source connectivity fails [43]. More importantly, Starlink and other LEO networks are being considered for nation-wide or even multi-national efforts to improve resilience in cases of terrestrial infrastructure failure. For example, HEIST is a NATO project to improve resilience of submarine cables through satellite-based communication [7, 38]. Other examples include Ghana, Taiwan, and Israel where Starlink is considered to be a failover network in case of wide-spread outages [19, 36, 76]. These examples reflect a growing willingness of governments to proactively incorporate LEO networks into national resilience strategies.

In this paper, we focus on scenarios where satellite infrastructure is considered proactively and systematically as failover at the national level. Our focus is motivated by the negotiation power that national governments have with satellite operators. Moreover, such plans will have far reaching effects when it comes to the funding of satellite networking infrastructures and the expectations of citizens regarding the resilience of their infrastructure.

LEO network regulation. A communication satellite has to comply with many regulations to limit its interference with other space-borne and terrestrial communication systems. Each nation regulates its own spectrum, licensing spectrum access to satellite operators [39]. The licensing process includes filing with the International Telecommunication Union (ITU), which tracks technical and operational parameters of communication satellites, including their transmit power, transmit beam contours, receiver sensitivity, and orbital parameters [40]. In contrast to cellular operators, satellite operators do not license exclusive access to bands and instead share these bands [22, 40]. Therefore, the capacity of failover LEO networks depends not only on decisions made by operators but also on government regulations on how RF spectrum is used, especially in emergency scenarios.





Figure 1: LEO network infrastructure. Radio-based Ground-Satellite Links (GSLs) connect satellites to gateways and terminals. Laser-based Inter-Satellite Links (ISLs) connect satellites to each other.



(a) A satellite can cover multiple cells using different beams. (b) The area of a cell is a small fraction of the total coverage area.

Figure 2: Illustration of RF resource allocation. Although a single satellite can cover multiple cells simultaneously, it can only cover a small fraction of cells within its coverage area due to the limited number of beams it can support.

2.2 Satellite Networks Technical Primer

A LEO satellite network, or a constellation, comprises thousands of satellites orbiting the Earth at altitudes in the range of 200–1600 km [20]. Satellites are placed in a number of shells, each consisting of multiple orbits (or orbital planes) at specific altitudes. Orbital planes in a shell are equally spaced and are characterized by their altitude (the height above sea level), and their inclination angle (the angle at which they intersect the equator). An inclination angle of 90° refers to a polar orbit. However, most of the current constellations have smaller inclination angles to provide greater coverage to densely populated areas [82].

Network customers use their terminals to access the satellite network through radio-based Ground-Satellite Links (GSLs). The network accesses the Internet through gateways that communicate with satellites through radio-based GSLs. A ground station only communicates with satellites that are visible above a certain *elevation angle* above the horizon, limiting the time traveled by the wave in the Earth’s atmosphere to ensure the quality of the link. There are two modes of communication that rely on satellites. First, the “bent pipe” scheme, where traffic goes through a single satellite hop before going back through a ground station. Second, data can travel through multiple satellite hops. Satellites communicate with each other through laser-based Inter-Satellite Links (ISLs). Figure 1 illustrates the components of the network.

RF Resource Management. There are two main resources in a satellite network: RF bandwidth used for communication between satellites and ground stations (i.e., user terminals and gateways) and

ISL bandwidth used for inter-satellite communication. To maximize RF bandwidth utilization, it is divided in time, space, and frequency. In particular, the coverage area of a satellite is divided into cells and the bandwidth is divided into channels. Each satellite uses highly directional beams to communicate with terminals in a specific subset of cells within its coverage area, with each beam using a particular channel. The radius of a cell is typically within a few tens of kilometers, while the total coverage area of a satellite has a diameter measured in hundreds of kilometers. The smaller the cells, the more efficiently RF capacity can be utilized. However, smaller cells require more satellites to provide full coverage and the technology to illuminate small cells with very narrow beams. In our analysis, countries with smaller areas would greatly benefit from using smaller cells to improve RF efficiency. On the other hand, countries with large areas are typically bottlenecked by the number of satellites that cover them. Satellites operators further improve RF efficiency by employing multiple directional antennas to enable frequency reuse (i.e., a satellite can transmit multiple beams on the same frequency). Figure 2 visualizes beam allocation and the scale of cells compared to the total coverage area. ITU and local regulations limit the power per beam and the contours of the coverage area of each beam [22, 40].

A network operator must decide how many beams to allocate per cell to maximize RF bandwidth utilization while providing service to all its customers. Beam allocation is constrained by potential interference between beams. In particular, two beams that use the same channel have to be physically separated or use a different polarization. When a satellite reuses a frequency, the beams using the same frequency must be angularly separated.¹ The beam allocation problem is known to be NP-hard [71]. Note that communication between satellites and user terminals typically employs a different frequency band than communication between satellites and gateways (e.g., OneWeb and Starlink). Moreover, gateways have a sophisticated setup (e.g., larger antennas and better receivers) to enable communication with multiple satellites while maximizing utilization of the RF spectrum. The utilization of ISLs is dictated by the traffic engineering policy employed by the network operator.

2.3 Related Work

Characterization of the performance of LEO satellite networks have received extensive attention. These include real-world measurements such as analyzing Starlink throughput and latency data [65], uncovering Starlink’s satellite-terminal scheduling [87], large-scale Starlink measurements to uncover network routing behavior [41], a multifaceted study evaluating Starlink’s ability for various tasks such as video conferencing and cloud gaming [66], studying Starlink behavior in vehicular and mobility scenarios [51], measuring the effectiveness of LEO networks for coastal maritime connectivity [25], measuring the asymmetry and variability in one-way delays for Starlink packets [35], and uncovering symbol allocations and modulation schemes via of burst characterization [96]. The Tiansuan constellation [91] presents the design for a research constellation to perform experiments on a real platform. LEOScope [30] introduces a global measurement infrastructure to perform real LEO network

measurements. The simulator designs presented in Hypatia [45] and xeoverse [44] help to explore routing and congestion control techniques for satellite networks. Other simulation-based work explores modeling available global bandwidth [24, 68, 69, 71], characterization of route variability [15, 16], vehicular satellite connectivity in a 3D foliage environment [53].

A common theme across these studies is their focus on per-user performance. Instead, we study the aggregate capacity available in a given country. Although some studies examine the global capacity of different LEO network constellations, their focus is on performance per customer and global coverage of the network for fixed demand [24, 68, 69, 71]. However, in cases of infrastructure failure, demand is subject to the needs of the populous and is primarily impacted by how governments distribute terminals that will be used by failover traffic. Moreover, we pay special attention to the policy implications for using LEO networks as a failover as it pertains to potential coordination between governments and network operators. Finally, our methodology is particularly unique in providing a clear framework for the capacity offered by satellite networks. It uses real submarine cable failures as a reference point for understanding the capabilities of LEO satellite networks. Additionally, we examine this framework in six different case studies.

3 Modeling LEO Failover Capacity

We assess the viability of LEO networks as a failover network in cases of national connectivity infrastructure failures.² We select real-world submarine cable failures as a reference to contextualize our analysis, providing an estimate of the demand that satellite networks could be expected to fulfill in similar scenarios. We assume that in such scenarios governments will undertake the deployment of a large number of user terminals, similar to what was done in Ukraine [89], and aggregate their bandwidth to offset capacity lost due to infrastructure failure. We analyze the failover capacity achieved by different terminal deployments under various RF spectrum allocation and traffic engineering policies.

Our analysis characterizes the impact of multiple deployment strategies. We refer to this problem as the terminal distribution problem and introduce several alternatives that yield different levels of performance (§ 3.1). Furthermore, performance depends on how satellite network operators allocate their resources (i.e., RF, ISL, and gateway capacity). No existing publicly available simulator captures all these aspects of a satellite network. Thus, we build a new simulator that analyzes the capacity of LEO networks under different terminal distribution and resource allocation policies while scaling to tens of thousands of satellites and hundreds of thousands of user terminals (§ 3.2). Our objective is to identify the key factors that determine the capacity of the satellite network to derive lessons learned and policies that can be followed for emergency planning and response. To this end, we implement and study multiple RF resource allocation and traffic engineering policies (§ 3.3).

3.1 The Terminal Distribution Problem

A unique problem faced by governments planning to leverage LEO satellites as failover is how to distribute their user terminals to

¹We enforce angular separation by preventing two beams that reuse the same frequency by the same satellite to illuminate neighboring cells.

²We limit the scope of this work to LEO satellite networks as a failover, and do not consider terrestrial alternatives such as cross-border or underground fiber cables, which may also contribute to resilience but are beyond the scope of this study.

serve their needs and maximize the aggregate network capacity. As discussed earlier, satellites allocate their capacity in beams that cover cells whose diameter is in the range of tens of kilometers. If a cell has too many terminals, the cell's RF capacity will get congested, limiting overall network capacity. Thus, we separate the terminal distribution problem into two components: 1) cross-cell distribution, deciding the aggregate number of terminals to be allocated to individual cells, and 2) local terminal distribution, deciding how terminals should be deployed within an individual cell. The intuition behind the two different scales of terminal distribution is the difference between the service capacity of individual terminals and the service capacity of satellites.

Cross-cell terminal distribution attempts to maximize the utilization of the RF capacity of satellites by spreading terminals between cells. On the other hand, local cell distribution attempts to meet the requirements of individual users or communities (e.g., individual buildings). Local terminal distribution accounts for user demand and the availability of infrastructure to deliver network capacity from terminals to individual users (e.g., WiFi or 5G). The local terminal distribution problem resembles cellular network planning [67, 75, 92], where backhauling is done through satellites instead of other forms of terrestrial networks. In this paper, we are concerned with the aggregate capacity of the satellite network, not the fine-grained planning of failover network design. Thus, we focus on the cross-cell terminal distribution problem.

The cross-cell terminal distribution problem exhibits a fundamental tradeoff between network utilization and deployment practicality. The number of user terminals in a cell plays a major role in determining the achieved capacity in that cell. However, the aggregate capacity provided to a single cell in the LEO network is capped by the RF bandwidth licensed to the network. On one hand, a government needs to deploy resources where they are needed, with more terminals placed in high density areas. On the other hand, the more terminals deployed to a particular location, the less the available RF bandwidth to individual terminals.

We assume that a government assigns each cell a priority level, reflecting needs, say based on population density, infrastructure criticality, or national security considerations. We assume governments will ensure a speedy deployment of terminals by pre-deploying them at regional depots or by pre-installing the critical infrastructure to be activated during the emergency. The assignment of terminals solely on the basis of priority sacrifices the aggregate capacity of the network by creating contention for spectrum capacity in high-priority regions. Spreading terminals evenly over different cells can lead to terminals being assigned to uninhabited cells, making it impractical to leverage the added capacity. We design a configurable heuristic that enables a government to perform terminal distribution along that spectrum.

The proposed algorithm is Greedy, Capped, and Batched, referred to herein as GCB. GCB is greedy, prioritizing cells with higher priority. The cap sets a lower limit on the priority of cells that can receive terminals, ensuring that terminals are allocated only to cells that meet a predefined criterion (i.e., high-priority cells). GCB assigns terminals to cells in large batches where a batch of terminals is enough to saturate the RF capacity in a given cell. The combination of greedy and batched allocation ensures that higher priority cells maximize their RF utilization. GCB can also




Figure 3: Construction of CosmoSim graph. Cell channel nodes represent the same set of cells, where each cell is covered by a different set of channels.

be uncapped, distributing terminals to all cells while prioritizing higher-priority cells. Uncapped GCB maximizes network capacity by avoiding wasting RF resources but can potentially assign terminals to sparsely populated cells. A cap set to a high priority value limits the number of cells that receive terminals, potentially limiting the overall network capacity. Without loss of generality, we assume population density as our proxy for priority. For example, a cap of 10,000 ensures that only cells with a population of 10,000 or more are assigned terminals. We refer to such an approach as **GCB (cap x)**, where x is applied cap. The exact steps are shown in Algorithm 1 in Appendix B. Although GCB significantly improves performance, governments may still need to customize the algorithm, targeting their unique requirements and population distributions.

3.2 CosmoSim

We built CosmoSim, a simulator to model aggregate network capacity under a wide range of terminal distribution, RF allocation, and traffic engineering policies, while scaling to tens of thousands of satellites and hundreds of thousands of user terminals. Existing simulators typically focus on per-connection performance. Instead, CosmoSim is concerned with the aggregate capacity on the network. To that end, we formulate our simulation as a maximum-flow problem. In particular, our analysis aims to compute the maximum possible flow between source nodes (gateways) to destination nodes (user terminals) subject to different policies. Policies determine the connectivity of the graph, the valid routes that flows can use, and the available capacity at individual edges. CosmoSim is lightweight, programmable, and modularized, enabling the community to not only further study the aggregate capacity of satellite networks but also experiment with different resource allocation policies and network loads and configurations. CosmoSim's parameters are instantiated from values reported in prior empirical and analytical studies of Starlink, including measured beam frame structures, beam allocation strategies, and aggregate throughput characterizations [24, 66, 68, 69, 71].

We model a LEO satellite network as a graph with five types of nodes: gateway nodes, satellite nodes, cell nodes, a source node, and a sink node. RF allocation restricts each cell to a fixed number of channels. To improve the flexibility of representing RF allocation policies, we represent each cell as a set of channel nodes, with each channel node corresponding to an individual channel used at a given cell. The source and sink nodes represent the source and destination of all traffic traversing the network from gateways to




Figure 4: The four main stages of CosmoSim in orange. First two stages can run in parallel. White boxes represent programmable components and configurations.

cells. CosmoSim can also capture incumbent traffic on the network. To that end, we add the cells from the areas that comprise the incumbent traffic (e.g., cells from all the countries served by the operator, not just the specific country we study). Each satellite can connect to other satellites with a number of edges limited by the number of ISLs it supports. Edges between satellites and ground stations (i.e., gateways and terminals) are based on satellite visibility from a given gateway or a cell and the RF allocation policy. The weight assigned to an edge represents the capacity of the link represented by the edges. Figure 3 illustrates this graph.

CosmoSim constructs its graph from three configurable inputs: (i) constellation configuration consisting of all satellite orbital parameters as well as ISL design, (ii) gateway locations and capacities, and (iii) cell characteristics and RF constraints/allocation policies. This enables CosmoSim to adapt to any LEO network constellation as long as these inputs are updated accordingly with the CosmoSim pipeline requiring no changes.

A naive implementation of the model presented in Figure 3 would imply processing graphs with hundreds of thousands of nodes for each network configuration. This would require running the simulator for several hours for every configuration (e.g., number of terminals and policies employed). To address this issue, we implement several optimization steps to reduce the number of nodes in the graph and reuse precomputed data to accelerate simulations.

To reduce the number of nodes in the graph, a first step is to forego representing each terminal as an individual node in the graph. Since RF allocation is performed on a per-cell basis, we can represent the number of terminals in a cell as a label of the node of that cell in the graph. Terminals within a cell are evenly distributed across its active channel nodes (i.e., channel nodes connected to a satellite). The label helps determine the maximum possible throughput achieved at the cell based on the number of terminals at the cell and the capacity of an individual terminal. Specifically, the capacity of a channel used at a particular cell is constrained by the minimum of the capacity of the terminals using that channel and the channel’s maximum capacity. This optimization helps to reduce the number of nodes in the graph by 100×, reducing processing time to under an hour. More details on the performance gains from the CosmoSim optimizations are in Appendix H.

CosmoSim divides its operations into four sequential phases, where the output of each phase is reusable under different configurations of the following phases in the pipeline. In particular, it has the following phases (Figure 4):

Graph Generation takes in as input satellite configuration (i.e., a two-line element set), gateway locations, and the cell configuration in a country to generate all nodes in the graph. ISL edges are defined based on the assumed topology of the satellite network. On the other hand, GSL edges are defined solely by visibility constraints. Specifically, a channel node of a cell or a gateway is connected to a satellite if that satellite is visible from the cell or gateway. Clearly, this stage creates many invalid edges which will be eliminated by subsequent processing phases. It is important to note that the graph generation step, being the most computationally intensive, is performed only once per country.

Terminal Distribution (TD) implements the cross-cell terminal distribution algorithm. In particular, it assigns labels to individual cells, reflecting the number of terminals assigned to that cell. This phase is performed once per country, per TD policy.

RF Allocation prunes the graph by eliminating GSL edges based on the RF allocation policy and wireless interference constraints. Moreover, this phase distributes the terminals allocated to a cell evenly between the channel nodes of that cell that have not been pruned. By this stage, we know precisely the demand placed on each satellite based on the cells it covers. Thus, to further reduce the size of the graph, we remove all channel nodes and replace them with the aggregate demand value at individual satellites. The aggregate demand value is used as the weight of the edge connecting a satellite node to the sink node.

Traffic Engineering determines the possible routes that traffic can take while traversing through the network graph. Taking as input the network graph, the TE policy updates the weights assigned to different edges in the graph. In the cases where we are also considering the incumbent traffic, the TE policy defines the demands from the satellites serving the incumbent cells to the sink. This final graph is used to estimate the capacity using max-flow. Specifically, we use the `max_flow_min_cost` function from NetworkX [8] that performs max-flow while reducing the the number of hops between terminals and gateways (i.e., the of a path cost). We use max-flow solely as an optimistic upper bound to estimate the maximum possible failover capacity, not as a representation of realistic traffic engineering policies.

The source code and the data for CosmoSim is available at <https://github.com/GT-ANSR-Lab/CosmoSim>. Further, we have published a web interface for the simulator available on the Github page. The supported features and a screenshot of the web interface are included in Appendix G.

3.3 CosmoSim Resource Allocation Policies

Extensibility is a core objective of CosmoSim. Thus, we simplify the process of implementing different terminal distribution, RF allocation, and traffic engineering policies. Moreover, we bootstrap the CosmoSim policy library by implementing and studying the following policies.

Terminal Distribution Policies. We consider the GCB policy under different values of the priority cap. Moreover, we compare it with the distribution of terminals based primarily on population density, representing scenarios where the government attempts to optimize primarily for local service coverage, ignoring overall network capacity.

RF Allocation. Our objective is to better understand the fundamental bottlenecks for LEO network capacity. Thus, we employ a greedy algorithm for beam allocation, similar to an algorithm that was shown to provide good performance [95], compared to other more sophisticated algorithms [55]. The simplicity of the greedy approach is particularly beneficial for our use case because it allows us to easily understand allocations created by the algorithm. More complex algorithms (e.g., multi-staged algorithms [11, 55]) would have been more complicated to analyze.

Specifically, our greedy beam allocation strategy prioritizes allocating beams to cells based on their priority. In addition, we prioritize the usage of satellites equipped with ISLs to ensure connectivity and reduce congestion at nearby ground stations. Of all satellites equipped with ISLs, we prioritize the least utilized satellites. Finally, we ensure that the allocated beams obey interference rules. In particular, we capture intra-satellite beam interference, ensuring a separation angle between beams when a satellite reuses the same frequency for the two beams. Our current interference model does not capture inter-satellite interference or degraded performance due to weather and terrain conditions. Thus, after allocating a beam to a cell, we prune the GSL edges that can interfere with it. Moreover, a cell is removed from the list once its demand is satisfied. We iteratively repeat the above process until all demand is satisfied or we all beams are exhausted. Our RF allocation is shown in Algorithm 2 in Appendix B.

We consider two variations of this greedy algorithm, differing based on the way they prioritize cells. The first prevents cell starvation, ensuring that each cell with at least one terminal is allocated at least one beam. Specifically, it ensures max-min fairness between cells by invoking the above algorithm repeatedly in the following way. The algorithm creates separate lists for all cells with demand of one or more beams, two or more beams, three or more beams, etc. Each list is sorted based on the population density of the cells. The above algorithm is invoked repeatedly for each of the lists until all demand is satisfied or all beams are exhausted. This algorithm is our best estimate of operator behavior which will avoid starving any of their clients while attempting to prioritize cells based on their estimated population density. However, it can severely underutilize the RF capacity by allocating beams to cells with not enough terminals to fully utilize the beam capacity. The second variation of the algorithm greedily prioritizes cells strictly based on the number of terminals at each cell. This proportional allocation maximizes overall network capacity by ensuring that the allocated beams can be fully utilized.

We do not aim to exhaustively identify the best RF allocation heuristic for failover satellite networks among the many recent proposals [55, 70, 71, 95]. Instead, we compare RF allocation policies that are aware of terminal distribution decisions (i.e., greedy proportional allocation) and those who are oblivious to it (i.e., attempting to achieve max-min fairness between cells).

Traffic Engineering Policies. We consider two traffic engineering policies: Max-flow and hot potato. Max-flow represents the best possible scenario where the network operator spreads the failover traffic across its whole network, avoiding congested links and balancing loads between all gateways. Hot potato represents the worst case where a network operator attempts to minimize the footprint of failover traffic on its network by routing traffic only through the

gateways nearest to user terminals. The two policies allow us to study the magnitude of the impact of network operator decisions on achieved performance.

4 Selected countries and cable cuts impact

To understand the impact of submarine cable failures on a country or region’s international connectivity, we estimate the relative importance of each cable from that region’s perspective. Prior work shows that the frequency with which a cable appears in traceroute paths serves as a proxy for the traffic volume it carries [27, 78]. Building on this insight, we analyze RIPE Atlas traceroutes from probes within the affected countries, using built-in measurements [3, 4] that target a broad cross-section of routed prefixes. We collect measurements over a baseline window of one week to one month prior to known submarine cable failure events. We use *Calypso* [90] a framework for mapping traceroute paths to the submarine cables they traverse, identify mappable routes, and distinguish domestic from international links. For each case, we compute the fraction of submarine-bound traceroutes mapped to each cable and use this ratio as a proxy for its relative importance. We complement this with capacity data from Telegeography [10], allowing us to estimate the potential gap for satellite networks to fill as backup infrastructure.

We apply this approach to six case studies that represent diverse conditions for LEO satellite coverage. These include differences in: (1) available land area for satellite cells, (2) proximity to ground gateways, (3) satellite density (affected by latitude), and (4) population distribution, which limits practical coverage. By analyzing scenarios with varying combinations of these factors, we explore the feasibility of satellite backup under realistic constraints. Thus, our selection of the case studies does not imply that using LEO satellite networks as failover is the best approach for handling outages in the countries we study. For example, South Africa is a particularly interesting case study from our perspective due to its large landmass and lack of nearby Starlink gateways.

For each case, we identify relevant cable failures, estimate the role of the cable in regional traffic, and calculate capacity loss. We use these values to gauge the adequacy of satellite networks to compensate for lost capacity.

Tonga – January 2022. Tonga, an island nation in the South Pacific, is highly vulnerable to natural disasters and relies entirely on a single submarine link—the TONGA CABLE—for international connectivity. With a maximum capacity of 320 Gbps [10], this cable connects Nuku’alofa to Fiji and the global Internet. On January 15, 2022, a volcanic eruption at Hunga Tonga-Hunga Ha’apai triggered a tsunami that severed the cable 37 km offshore, resulting in a near-total communications blackout [64]. A similar failure in 2019 left the country offline for over 10 days. We analyzed one month of RIPE Atlas traceroutes from all probes in Tonga and found that 92% relied on submarine cables, with the TONGA CABLE ranked as the most critical path—underscoring the nation’s complete dependence on a single international link [18]. Details on the ranking method are in Appendix C. While submarine infrastructure development has been minimal in recent years [94], Tonga has improved resilience since the 2022 outage by partnering with multiple satellite providers—including SpaceX’s Starlink, Kacific, and SES—to provide alternative connectivity [37, 56, 80].

South Africa — August 2023. South Africa, the continent’s second-largest economy, hosts landing stations for ten international submarine cables, including SAT-3/WASC, SAFE, SEACOM, EASSy, and METISS. One month of RIPE Atlas traceroutes shows that 87% of outbound paths rely on submarine links, underscoring their central role in the country’s global connectivity. On August 6, 2023, an underwater landslide in the Congo Canyon disrupted two key west-coast cables—SAT-3/WASC and WACS—causing widespread service degradation across southern Africa [59] (an example of one Atlas probe displaying increased latency, jitter, and packet loss is included in Appendix E).

Both cables rank among the top three in our criticality analysis. Together, they provide 10,700 Gbps of lost capacity, with a potential maximum of 31,700 Gbps if fully upgraded [10], making them essential for regional and international transit.

Ghana — March 2024. Ghana, a key player in West Africa’s Internet connectivity, relies heavily on submarine cables for global access. Its Gulf of Guinea coastline hosts six major cables, including WACS, MAINONE, SAT-3/WASC, ACE, GLO-1, and 2AFRICA. On March 14, 2024, a suspected underwater rockslide off Côte d’Ivoire disrupted Ghana’s connectivity by taking four major cables offline: ACE, SAT-3/WASC, WACS, and MAINONE [81]. Analyzing one month of traceroutes from all probes, we found that 90% relied on submarine cables, highlighting their critical role. SAT-3/WASC ranked as the most essential, with ACE, WACS, and MAINONE also playing significant roles. These rankings reflect the high concentration of traffic on these routes, emphasizing the need to identify vulnerabilities and guide contingency planning. The four affected cables have a combined lit capacity of 57,170 Gbps, with a potential maximum of 83,700 Gbps if fully upgraded, according to TeleGeography.

Additional case studies are detailed in Appendix C.

5 LEO Network Failover Capacity Analysis

We study the different factors and policies that impact the failover capacity provided by LEO satellite networks, using Starlink as a representation of available satellite-based network capacity. Our analysis captures the capacity of Starlink deployment at the time of writing this paper. In addition, we compare the current available capacity with that available as Starlink expands, relying on FCC filings to estimate the configuration of future Starlink deployments [31]. Analyzing the expanded Starlink network is significant as it does not only reflect the value of Starlink growth but also the deployment of other constellations. In particular, all current constellations share the same spectrum [32, 48, 60]. Thus, expansions of Starlink or any other constellation will lead to competition between satellites for the same RF bandwidth.

Our findings (marked by $\mathcal{F}$) highlight critical bottlenecks in failover LEO networks. Based on our findings, we make policy recommendations to governments and operators (marked by $\mathcal{R}$). We also provide a complementary policy-focused analysis of our findings and recommendations [17]. Our recommendations aim to facilitate the attainment of maximum network capacity. Further, they highlight areas that require further study. In particular, we answer the following questions:

- How does terminal distribution affect available capacity? (§5.2)
- Should governments and network operators coordinate their resource allocation efforts? (§5.3)

Data Rates (Gbps)	Config A	Config B	Config C
Ku GSL beams (Users)	1.28 [72]	0.956 [70]	2.5
Ka GSL beams (GWs)	2.6 [72]	2.6 [72]	5
ISLs	100	100	200

Table 2: Different wireless configurations used in this paper.

- Will LEO network growth increase the available capacity for failover traffic? (§5.4)
- What’s the impact of sharing LEO network capacity with failover traffic? (§5.5)

5.1 Simulator Configuration

We use the configurations drawn from current and planned Starlink deployments, the only LEO satellite network with thousands of satellites and the most prominent failover network in many recent scenarios [21, 38, 43, 58, 74, 79].

Topology. Starlink has about 6,500 satellites deployed in space across five different shells [42]. All satellites, except those in shell 1 have three full-duplex ISLs with two connecting to satellites in the same orbital plane and one to a satellite in a neighboring orbital plane (more details in §D). The FCC has announced the approval of two more shells while deferring approval for another six [31]. For satellite orbital configuration, we use Celestrak [46] two-line element sets (TLEs) to derive the configurations but generate our own idealized version to facilitate the creation of ISL-based topology. We obtain the location of Starlink gateways from FCC filings (e.g., [2]). We use 198 gateways spread over 23 countries. For cell configuration, Starlink has been reported to divide its coverage region into hexagons, following a scheme similar to Uber’s H3 hierarchical spatial indexing system [54, 57]. We assign the population density for each cell using the Kontur.io dataset [49].

Wireless Link Data Rates. The data rates achieved by wireless GSLs depend on many technological (e.g., modulation and coding schemes) and environmental factors (e.g., weather). Instead of using a single value for the capacity of each type of GSL, we consider two different estimates made in prior work [24, 70, 72, 77]: Config A and Config B in Table 2. Moreover, we add an estimate of wireless capacities, assuming technological advancements in next-generation LEO satellites that double the capacity of all links (Config C). Our objective is to compensate for the inaccuracy of our simulation-based approach by providing an envelope on network performance that captures most plausible network behaviors. Our estimate of ISL capacity is based publicly available technology specifications from Starlink [9]. These GSL and ISL capacities represent the upper bound estimates as we are estimating an upper bound for the failover capacity. Real-world conditions such as weather and terrain may yield lower values.

We use Config A for most of our analysis, preferring the optimistic estimate of channel capacity when communicating with users. In particular, a single beam connecting a satellite to a cell for user communication has a capacity of 1.28 Gbps. Each satellite can communicate over eight such channels with a frequency reuse factor of four [24], resulting in a total per satellite capacity of 40.96 Gbps for communication with user terminals. On the other hand, each Ka band GSL beam connecting from the gateways to the satellites has data rate 2.6 Gbps. With eight active gateways per site, total gateway site capacity reaches 166.4 Gbps which is the capacity




Figure 5: Estimated capacity for various terminal distribution strategies indicates that distributing terminals based solely on population density can reduce network capacity by causing RF contention in densely populated areas and not allocating enough terminals to other areas with RF capacity.

of all the source-gateway links in our graph. Our analysis focuses on user downlink capacity, since it dominates aggregate demand in failover scenarios. Uplink capacity differs due to spectrum allocation asymmetry. For instance, the user uplink capacity for Starlink will be one quarter that of the downlink due to the differences in the spectrum allocation. While our analysis does not include any gateway failures to limit the number of changing variables, CosmoSim design can be easily extended to incorporate those by simply omitting the failed gateways from the list of gateways.

Default policies. Unless otherwise stated, we use the following policies by default. For terminal distribution, we use GCB. For RF allocation, we use the greedy policy that allocates capacity purely based on the number of deployed terminals by the government, potentially leading to the starvation of other users while maximizing failover capacity. For traffic engineering, we use max-flow. These default policies favor the maximization of failover capacity. However, we investigate the impact of other policies as well.

5.2 Impact of Terminal Distribution

Terminal distribution decisions must balance maximizing spectrum utilization, by spreading terminals between coverage cells, and ensuring that terminals are deployed where they are needed, potentially leading to resource contention between them. In this section, we focus on assessing the maximum achievable capacity, assuming that there is no other traffic on the network and that all cells can potentially be used by emergency responders, including completely uninhabited remote areas.

Figure 5 reports the total capacity achieved in each of the case studies under different number of terminals and terminal distribution strategies. Clearly, increasing the number of terminals increases aggregate capacity. Moreover, the GCB algorithm improves failover capacity compared to distributing terminals based on population density, especially when the number of terminals is small relative to the number of cells. When the number of terminals is limited and they get allocated based on population density, *areas with very high population density end up getting most of the terminals, creating an RF bottleneck* ($\mathcal{F}$ 1). For example, in case of Lithuania,

the capacity jumps by over 3 times when terminals are deployed across cells using GCB instead of based on their population density. In such scenarios, more terminals are needed so that sparsely populated cells can receive enough terminals to consume the resources of beams allocated to them. *Therefore, terminal distribution needs to balance population density and RF availability* ($\mathcal{F}$ 1).

Capping the minimum population density of cells that receive terminals can limit available capacity. However, we find that, for populous countries, the impact is most significant when the limit is 100k. A limit of 100k can reduce the maximum achievable capacity of a country to 2.5-65% of the maximum achievable capacity. *The impact of focusing on densely populated area depends primarily on how a country's population is spread over its landmass.* ($\mathcal{F}$ 2). We consider the impact moderate if it can achieve 50% or more of maximum capacity (e.g., Britain and Tonga). We consider the impact significant otherwise. Although South Africa and Great Britain have comparable populations with 60.41 and 68.35 million, respectively, the impact of the cap is much more severe in the case of South Africa (72% reduction in capacity) compared to Britain (40% reduction). The reason is that Britain's population is more spread out over its landmass, leading to more cells with over 100k people compared to South Africa. Obviously, Tonga does not even have a single cell with that 100k people, making it infeasible to apply that limit. A 1k cap achieves the maximum possible throughput with a comparable number of terminals to a cap of zero.

Increasing the number of terminals yields diminishing returns in large deployments. For small deployments, adding terminals yields a linear increase in aggregate capacity because each added terminal makes use of some beam capacity. However, large deployments, starting with 10-20k terminals and beyond, yield diminishing returns with additional terminals. These diminishing returns are not an outcome of network congestion, but rather due to some cells receiving more terminals than the capacity of all the beams assigned to them. Thus, *deployments with a large number of user terminals exhaust the capacity of all satellites visible from the country* ($\mathcal{F}$ 3).

While bandwidth can always be viewed as the only bottleneck, there are two key factors that play a role in limiting the bandwidth




Figure 6: CDF of GSL and ISL utilization for South Africa when 200k terminals with the GCB policy with no cap.




Figure 7: Utilization for different gateways for South Africa when 200k terminals with the GCB policy with no cap.

available for a given country: the land area of the country and the number of satellites covering the country. Countries with a small land area, like Tonga (4 cells) and Haiti (136 cells), can have all their cells receiving 8 beams, saturating the RF capacity of each cell when each cell is allocated 200 terminals. Any additional terminals, or satellites, will not increase the capacity of the satellite network. Therefore, *the remedy for smaller countries is either improve their RF efficiency using smaller cells and narrow beams or to allocate more RF bandwidth to satellite networks* ($\mathcal{R}$ 2). In contrast, the larger the land area of a country, the more satellites it will need so that each cell receives 8 beams. Thus, countries with large landmasses are bottlenecked on the number of deployed satellites. For example, South Africa, with 4545 cells barely receives one beam per cell. The bottleneck is exacerbated because South Africa cannot utilize any of the satellites in the first shell of Starlink, as “bent-pipe” satellites require the presence of nearby gateways, which South Africa does not have.³

An important concern when studying aggregate network capacity is congestion. To examine the impact of congestion, we focus on the case study of South Africa. In particular, South Africa achieves the highest aggregate capacity due to its large landmass. Moreover, South Africa does not have any gateways nearby, forcing it to rely primarily on ISLs, increasing the chances of ISL congestion. First, we examine the levels of GSL utilization (Figure 6a). We find that only 35% of gateways are utilized with just 15% of them exhausting more than 50% of their capacity. Our traffic engineering approach attempts to minimize the number of hops used, even when max-flow is employed. Thus, gateways located closer to South Africa,

³The nearest gateway to South Africa is in Ghana, requiring the use of ISLs to reach it.




Figure 8: Utilization and number of hops from the terminals for ISLs for the South Africa case when 200k terminals are distributed using the GCB policy with no caps.

like those in Ghana, Chile, and Australia, achieve higher utilization (Figure 7). Next, we examine ISL utilization (Figure 6b). Only 6% of ISLs are ever used in the South Africa case study, with less than 1% of ISLs achieving near 100% utilization. To better understand ISL utilization, we focus on utilized ISLs (919 out of 14k) and plot a 2D histogram as a function of ISL utilization and hop count from user terminals (Figure 8). As expected, most ISLs are barely utilized, including those close to user terminals. Highly utilized ISLs belong to two categories: 1) those close to user terminals, potentially carrying the traffic of its satellite’s GSL as well as other traffic of nearby satellites, and 2) those close to gateways, carrying the aggregate traffic of the country. Countries with small landmasses exhibit even fewer ISLs and gateways with high utilization. Thus, *we conclude that congestion is a nonissue, if failover traffic is prioritized or provided exclusive access to the LEO network* ($\mathcal{F}$ 4).

5.3 Impact of Network Operator Policies

Impact of spectrum allocation decision. Recall that in normal scenarios beam allocation is done independently of the terminal distribution. In particular, the LEO network operator can require having at least one beam allocated to very sparsely populated cells, with remaining beams distributed based on population density or customer presence. Starlink targets rural communities and recreational travelers operating in remote areas. For example, Britain can have many Starlink users using its network in sparsely populated regions of Scotland. Moreover, its customers include airlines that fly their aircrafts over sparsely populated areas. Thus, *the number of beams allocated to failover traffic, and hence its capacity can be affected dramatically by the spectrum allocation policy* ($\mathcal{F}$ 5). Figure 9 shows the impact of coordinating beam allocation to maximize the capacity available to failover traffic, compared to spreading beams uniformly over a sparsely populated area of Britain. Lack of coordination can reduce the available capacity by 42% when using 100k terminals. The scarcity of the RF spectrum requires careful management of the resource. *We argue that countries should stipulate in their laws and spectrum licenses how spectrum should be used by satellite networks in cases of national emergency* ($\mathcal{R}$ 3). The spectrum should be allocated in the best way to meet the needs of the country. We believe that such an approach is feasible given the leverage that governments have with the contract sizes currently planned for such infrastructure [14, 29, 52]. Combining $\mathcal{R}$ 1 and




Figure 9: Estimated capacity of Great Britain to distribute terminals by population density, emphasizing the advantages of coordinated spectrum allocation and terminal distribution.




Figure 10: Comparing failover capacity while using max flow and hot potato routing for uniformly distributed terminals capped at 100k population cells.




Figure 11: Ghana's failover capacity with different gateway restrictions shows that reasonable data sovereignty constraints (e.g., All except Nigeria) do not impact capacity.

$\mathcal{R}$ 3 improves the capacity of the network at 200k terminals by 1.7–1.8 $\times$ for countries with large landmasses. Alternatively, combining the two recommendations can help reduce the number of terminals needed to achieve high capacity. For example, combining $\mathcal{R}$ 1 and $\mathcal{R}$ 3 achieves 1.3 $\times$ more capacity with 4 $\times$ less terminals in the case of South Africa, compared to using max-min fair RF allocation and terminal distribution based solely on population density.

Beams should be allocated to cells based on their utility from the perspective of the government facing infrastructure failure. In particular, a government assigns each cell a priority level, say based on population density, infrastructure criticality, or national security considerations. This problem can be formulated as finding a configuration of terminal distribution and beam distribution to cells to maximize the overall bandwidth utilization while minimizing interference between beams. Exploring such algorithms to increase failover capacity is left for future work. The tight coupling of terminal distribution and spectrum allocation might be untenable in practice. A simpler alternative could be to clearly mark the terminals used for failover traffic (e.g., with their MAC address). The network operator then provisions beams exclusively for those terminals. Note that this is done only in cases of national emergency by dedicating the spectrum to emergency efforts.

Impact of traffic engineering policies. Figure 10 shows the normalized capacity achieved using max flow and hot potato policies. Hot potato limits the number of gateways that a country can use to one or two, based on the hop count between satellites serving the country and gateways. Countries with bottlenecks on the number of cells they can use, due to smaller land mass or population concentration in a small number of cells, are less affected by hot potato routing. On the other hand, Britain is significantly affected because it has significantly higher traffic generated from and near London that gets routed to the Villeneuve-d'Ornon gateway in France leading to congestion along those routes. Hot potato routing, which optimizes primarily for latency and network utilization, leads to abysmal results. Thus, it is important that traffic is spread across many gateways, prompting data sovereignty concerns.

Impact of data sovereignty rules. The potential sensitivity of failover traffic can require it to go only through gateways deployed in friendly nations. We study the effect of a gateway mask that a country can use to list its approved gateways and then use the max-flow TE policy to direct traffic to those approved gateways. We focus on the case study of Ghana, the only case study where only one neighboring country, Nigeria, having Starlink gateways and all




Figure 12: Even with uniform terminal deployment and no caps, failover capacity plateaus due to spectrum exhaustion at or before 34k satellites as the LEO network size increases.




Figure 13: Impact of capped terminal distribution on South Africa's estimated capacity with an expanding LEO network, highlighting how usable capacity quickly plateaus due to RF bottlenecks in densely populated areas.

other gateways farther away. For all other case studies, either all gateways are far or several neighboring countries have gateways. Figure 11 shows normalized capacity when only Nigeria's gateways are used and when all other gateways, except for Nigeria's are used. Clearly, using very strict constraints, like with hot potato routing, significantly impacts achieved capacity. However, none of our case studies were bottlenecked on aggregate gateway capacity. Thus, as long as the data sovereignty constraints are not too restrictive, they will have no impact on aggregate capacity ($\mathcal{F}$ 6).

5.4 Impact of LEO network growth

As discussed earlier, Starlink and other LEO networks plan to dramatically grow their capacities over the next few years. In addition, the main bottleneck for countries with large land areas is the number of satellites that cover the country. Thus, we explore the impact of increasing the capacity of the LEO network. We use three configurations – (i) fully deployed Starlink shells S1 - S4 and partially deployed S5 (6,400 satellites), (ii) all the deployed and approved Starlink shells S1 - S7 (14k satellites), and (iii) all the deployed, approved, and proposed shells S1 - S13 (34k satellites). Additionally, we increase the number of terminals to determine if the increased capacity can be fully utilized by the nations. Terminals are distributed uniformly with a cap of zero on cell population.

Figure 12 presents the capacity matrix achieved by nearly doubling the number of satellites on the y-axis and doubling the number of terminals on the x-axis. As expected, both Tonga and Haiti see minimal gains with increasing number of satellites, since all cells receive the maximum number of beams even with fewer satellites. Lithuania approaches 8 beams per cell with 14k satellites. Conversely, Ghana, Britain and South Africa nearly double their capacity each time the number of satellites is doubled. Furthermore, these nations require even more terminals to realize the maximum capacity when the entire constellation is deployed with 34k satellites. Crucially, we observe that *all the case studies plateau at or before 34k satellites, having fully exhausted their spectrum* (§ 7). Adding more satellites would not increase the capacity for any of these nations, as they are bottlenecked by their available spectrum.

Adding a cap to the population density of cells that receive terminals further limits the value of increasing the size of the LEO satellite network. Figure 13 shows the impact of increasing the number of satellites and terminals on aggregate capacity when with different caps in the case of South Africa. The maximum capacity achievable in South Africa is 28 Tbps when the number of satellites is 34,000 and a cap of zero is applied to the minimum population density of serviced cells. In contrast, applying a cap of 100k reduces available capacity to 1 Tbps, regardless of the number of satellites in the constellation. A cap of 10k improves the maximum achievable capacity up to 9 Tbps. Thus, *even for larger countries, the easily-usable capacity nears its plateau with the main bottleneck being RF contention within densely populated areas* (§ 8). This phenomenon could already be observed in some US cities, where Starlink is no longer accepting new customers due to its network being at capacity [86]. Overcoming this problem will require further innovation in the management of RF resources, including the use of narrower beams. Moreover, governments can license more RF capacity to satellite operators as demand grows.

Impact of wireless channel capacity. The above results were generated using Config A for wireless channel capacity (Table 2). We reproduced the results in Fig. 5 using Config B and C. We observe similar trends. However, Config B reduces network capacity by 12.5-28%. On the other hand, Config C assumed better data rates improves network capacity by 1.8-5.5 $\times$ while requiring up to 2.5 $\times$ more terminals to achieve the maximum possible capacity. However, even under such an extremely optimistic outlook for LEO networks, they can only support a fraction of submarine cable capacity (e.g.,

32% of the capacity of a single cable in South Africa). The detailed results are presented in Appendix A.

5.5 Impact of Sharing the Network

Our analysis thus far assumes that failover traffic is the only user of the network. Now we examine the impact of other users, referred to as “Incumbent Traffic.” We assume that incumbent demand is uniform across all satellites. We only consider demand originating from the countries where Starlink currently offers its service. We make this simplifying assumption as there is no publicly available ISP-level usage data for Starlink. While this traffic matrix does not capture any local traffic heterogeneity, it is more representative of Starlink operations today compared to the ones used in prior work [12, 16, 45] that focused solely on traffic between the top 100 most populated cities.⁴ We focus on terminal distribution strategies that maximize capacity for different caps on the population size of serviced terminals.

Relative priority between failover traffic and incumbent traffic can impact the performance and cost of the failover solution. Prioritizing failover traffic to extract maximum capacity can come at a premium while potentially impacting the performance of incumbents. Examining the impact of high-priority failover traffic on low-priority incumbent traffic in aggregate is not very informative given the geographical spread of incumbents and their relatively higher demand. Thus, we study the magnitude of the impact on incumbents as a function of its spatial spread. Figure 14 is a heatmap demonstrating normalized incumbent capacity per cell across Europe and Asia. The heatmap value is assigned per cell as the average of the percentage of incumbent traffic served by the satellites visible to the cell. We assume 8.4 Gbps incumbent demand for every satellite serving incumbent traffic. *The high failover demand of Britain has far-reaching impact on Starlink’s global coverage* (§ 9). This affects most of Europe with some of the neighboring areas only receiving 50% of normal capacity. Further, this impact is even felt in Iran and Mongolia that rely on the gateways in Europe. South Africa has a similar impact on global coverage. However, other case studies have localized or no impact. We discuss the local impact of Ghana’s traffic in Appendix F. Conversely, *a LEO network with heavy incumbent traffic significantly impacts the failover capacity, to the extent of making it unusable* (§ 10). We discuss these findings in Appendix F

6 Discussion and Limitations

Table 1 summarizes our findings for the six case studies. The two nations that can reasonably compensate for lost submarine cable capacity are Haiti and Lithuania, largely due to relatively low capacity of their existing infrastructure. In the four other scenarios, LEO networks can compensate for a very small portion of lost capacity. Moreover, achieving that maximum theoretical capacity assumes that terminals can be deployed in remote regions and that the network is dedicating its resources to the country in the case study. Some case studies require 200,000 terminals to achieve that maximum capacity. We note that the deployment size is not a concern. For example, Ukraine received over forty thousand terminals to help offset failures in its infrastructure during the war [89]. The

⁴Note that most of the 100 most populated cities do not have gateways. Thus, pairwise traffic matrices between specific cities is not at all representative of Starlink operations.




Figure 14: Great Britain’s failover traffic significantly affects deprioritized incumbent traffic across most of Europe and Asia. Neighboring countries lose up to 50% of their LEO network capacity, highlighting the potential need for international coordination in cases of infrastructure failure.

distribution of terminals based on population density is a more realistic approach, as it provides capacity in regions where it can be used. However, it can yield worse results, especially when the number of terminals is limited. Attempting to distribute terminals more evenly between cells, even when capping the minimum population size of serviced cells, can significantly impact performance.

Limitations. This paper provides an optimistic approximation of the available capacity on a LEO network. In particular, we assume that all wireless and optical links operate at their maximum capacity. Moreover, we assume that any terminal deployed by a government in an emergency scenario will be fully utilized. These assumptions can be violated with poor weather conditions and wireless interference (due to other usage), limiting the capacity of wireless links. We also assume perfect line-of-sight for all the terminals and do not consider impacts due to terrain that may impact the wireless capacity. Our interference model only accounts for intra-satellite interference and does not capture interference due to other constellations. In addition, logistical challenges can limit the use of terminals. We choose to err on the side of optimism to provide an upper bound on the best achievable performance under the different conditions that we study. We assume that gateways are always available and not affected by the submarine cable outages. Our formulation of the capacity using a static max-flow formulation serves only as an upper bound of the failover capacity. We show that despite this optimistic approach, the achieved capacity of satellite networks is typically a very small fraction of that achievable by terrestrial and submarine infrastructure.

We focus our study on a small subset of resource allocation policies for traffic engineering, spectrum allocation, and terminal distribution. Our chosen policies reflect extreme approaches, with more nuanced policies likely to provide compromises between the extreme scenarios we select. For example, we compare max flow and hot potato traffic engineering, proving the best possible throughput with the first and extreme case of gateway congestion in the latter. Our goal is not to declare a winner between the policies, but rather to show how choices made by governments and network operators can impact network performance.

We recognize that governments may pursue different policy objectives in their approach to emergency connectivity. Some may prioritize maximizing total national failover capacity, while others may emphasize broad geographic coverage to ensure connectivity

across critical sites and essential services. The feasibility of government and LEO operator coordination, which underpins several of our recommendations, depends on commercial and political realities. Encouragingly, recent contracts with Ghana [36], Israel [76], and NATO [7] show that such agreements are becoming increasingly plausible, and we hope our work provides a framework in aiding those discussions.

7 Conclusion

LEO satellite networks, and Starlink in particular, are increasingly seen as a fallback for national-scale outages, especially during submarine cable failures. We evaluated their viability through six case studies in various regions. In most cases, LEO networks could replace only a small share (0.9%–14.7%) of lost capacity. Our findings highlight key factors of available LEO capacity including country size, terminal density, spectrum policies, and traffic engineering. Optimizing terminal placement and spectrum allocation can improve capacity, but even under ideal conditions, LEO networks remain a supplement, not a substitute, for submarine links. Beyond technical considerations, our results highlight the need for coordination between governments and satellite operators. Strategic agreements on failover capacity, emergency integration, and traffic prioritization are essential, improving capacity by up to 1.8× while requiring the same number of terminals or even fewer. However, prioritizing failover traffic can have unintended global effects on network coverage. Future work should explore dynamic spectrum use, satellite network design, and economic models for sustained emergency coverage.

Acknowledgments

We thank the anonymous reviewers and our anonymous shepherd for their constructive feedback that helped us improve this paper. This project was partially supported by NSF grants CNS-2345827 and CNS-2107392.

References

- [1] 2005. fibralink. Available at <https://www.semanticnet.net/subseacable/IZc1wMFSrEyWcbwhaj1/fibralink>. Accessed on September 29th, 2025.
- [2] 2019. FCC Public Notice re: Satellite Radio Applications Accepted For Filing. Accessed on September 29th, 2025.
- [3] 2024. RIPE Built-in Measurement 5051. Available at <https://atlas.ripe.net/measurements/5051>. Accessed on September 29th, 2025.
- [4] 2024. RIPE Built-in Measurement 5151. Available at <https://atlas.ripe.net/measurements/5151>. Accessed on September 29th, 2025.
- [5] 2024. SpaceX seeks to modify its existing authorization to change the orbital configuration and certain radio frequency characteristics of the authorized shells of the Gen2 system. Available at https://licensing.fcc.gov/cgi-bin/ws.exe/prod/ib/forms/reports/swr031b.htm?q_set=V_SITE_ANTENNA_FREQ.file_numberC/File+Number/%3D/SATMOD2024101100224&prepare=&column=V_SITE_ANTENNA_FREQ.file_numberC/File+Number. Accessed on September 29th, 2025.
- [6] 2024. Submarine Cable Map. Available at <https://www.submarinecablemap.com/country/united-kingdom>. Accessed on September 29th, 2025.
- [7] 2025. Hybrid Space-Submarine Architecture Ensuring Infosec of Telecommunications (HEIST) - A NATO Project. Available at <https://natoheist.org/>. Accessed on September 29th, 2025.
- [8] 2025. max_flow_min_cost — NetworkX 3.4.2 documentation. Available at https://networkx.org/documentation/stable/reference/algorithms/generated/networkx.algorithms.flow.max_flow_min_cost.html. Accessed on September 29th, 2025.
- [9] 2025. Starlink - Technology. Available at <https://www.starlink.com/am/technology>. Accessed on September 29th, 2025.
- [10] 2025. Telegeography. Available at <https://www2.telegeography.com/>. Accessed on September 29th, 2025.
- [11] Alexis I Aravanis, Bhavani Shankar MR, Pantelis-Daniel Arapoglou, Gregoire Danoy, Panayotis G Cottis, and Björn Ottersten. 2015. Power Allocation in

- Multibeam Satellite Systems: A Two-Stage Multi-Objective Optimization. *IEEE Transactions on Wireless Communications* 14, 6 (2015), 3171–3182.
- [12] Suvam Basak, Amitangshu Pal, and Debopam Bhattacharjee. 2025. {LEOCraft}: Towards Designing Performant {LEO} Networks. In *2025 USENIX Annual Technical Conference (USENIX ATC 25)*. 789–813.
 - [13] David Belson. 2024. Resilient Internet connectivity in Europe mitigates impact from multiple cable cuts. Available at <https://blog.cloudflare.com/fr-fr/resilient-internet-connectivity-baltic-cable-cuts/>. Accessed on September 29th, 2025.
 - [14] Ari Bertenthal. 2024. Ontario Awards \$100 Million Internet Contract to Starlink. Available at <https://broadbandbreakfast.com/ontario-awards-100-million-internet-contract-to-starlink/>. Accessed on September 29th, 2025.
 - [15] Vaibhav Bhosale, Ketan Bhardwaj, and Ahmed Saeed. 2023. Astrolabe: modeling RTT variability in LEO networks. In *Proceedings of the 1st ACM Workshop on LEO Networking and Communication*. 7–12.
 - [16] Vaibhav Bhosale, Ahmed Saeed, Ketan Bhardwaj, and Ada Gavrilovska. 2023. A characterization of route variability in leo satellite networks. In *International Conference on Passive and Active Network Measurement*. Springer, 313–342.
 - [17] Vaibhav Bhosale, Ying Zhang, Sameer Kapoor, Robin Kim, Miguel Schlicht, Muskaan Gupta, Ekaterina Tumanova, Zachary Bischof, Fabian Bustamante, Alberto Dainotti, and Ahmed Saeed. 2025. Are LEO Networks the Future of National Emergency Failover? – A Quantitative Study and Policy Blueprint. *SSRN Electronic Journal* (July 2025). <https://doi.org/10.2139/ssrn.5337566> Available at SSRN: <https://ssrn.com/abstract=5337566>.
 - [18] Kristin Carlson. 2022. A Volcanic Eruption Cut Tonga's Only Undersea Cable Connection. Available at <https://web.archive.org/web/20250212014745/https://blog.telegeography.com/telecom-headlines-january-2022>. Accessed Web Archive link on September 29th, 2025.
 - [19] Eric Cheung. 2024. Developing Taiwan's own 'Starlink' crucial for island-wide emergency, space agency says. (Accessed on September 29th, 2025).
 - [20] Christopher S. Allen and Martina Giraudo and Claudio Moratto and Nobuyasu Yamaguchi. 2018. Chapter 4 - Spaceflight environment. In *Space Safety and Human Performance*, Tommaso Sgobba and Barbara Kanki and Jean-François Clervoy and Gro Mjeldheim Sandal (Ed.). 87–138.
 - [21] Kevin Collier. 2022. Starlink internet becomes a lifeline for Ukrainians. Available at <https://www.nbcnews.com/tech/security/elon-musk-starlink-internet-becomes-lifeline-ukrainians-rcna25360>. Accessed on September 29th, 2025.
 - [22] Federal Communications Commission. 2020. Further Streamlining Part 25 Rules Governing Satellite Services (IB Docket No. 18-314).
 - [23] Matthew Davies. 2024. Britain's Undersea Cables Vulnerable in the Face of Security Threats, Report Warns. Available at <https://subtelforum.com/britains-undersea-cables-vulnerable/>. Accessed on September 29th, 2025.
 - [24] Inigo Del Portillo, Bruce G Cameron, and Edward F Crawley. 2019. A technical comparison of three low earth orbit satellite constellation systems to provide global broadband. *Acta astronautica* 159 (2019), 123–135.
 - [25] Birkan Denizer, Nils Dohse, and Olaf Landsiedel. 2025. FjordLink: Comparison of Starlink and 5G Networks for Teleoperated Vessel Control. In *Proceedings of the 2025 3rd Workshop on LEO Networking and Communication*. 8–14.
 - [26] John S Downs, Bruce Barbour, Alexander Kedrowitsch, Deven Mhadgut, Suryansh Aryan, and Samantha P Kenyon. 2025. Space Network (SpaceNet) Testbed-Development of a Multi-Functional Testbed for Simulating Space Communication Networks. In *AIAA SCITECH 2025 Forum*. 2716.
 - [27] Ramakrishnan Durairajan, Paul Barford, Joel Sommers, and Walter Willinger. 2015. InterTubes: A Study of the US Long-haul Fiber-optic Infrastructure. In *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication* (London, United Kingdom) (SIGCOMM '15). Association for Computing Machinery, New York, NY, USA, 565–578. <https://doi.org/10.1145/2785956.2787499>
 - [28] The Global Economy. 2016. Lithuania: International Internet Bandwidth per User. Available at https://www.theglobaleconomy.com/Lithuania/Internet_bandwidth/. Accessed on September 29th, 2025.
 - [29] Sandra Erwin. 2024. Starlink set to hit \$11.8 billion revenue in 2025, boosted by military contracts. Available at <https://spacenews.com/starlink-set-to-hit-11-8-billion-revenue-in-2025-boosted-by-military-contracts/>. Accessed on September 29th, 2025.
 - [30] Saeed Fadaei, Shubham Tiwari, Aryan Taneja, Saksham Bhushan, Mohamed Kassem, Aravindh Raman, Debopam Bhattacharjee, Lili Qiu, Alan Woodward, and Nishanth Sastry. 2025. Leoscope: Building a global testbed for low-earth orbit satellite networks. *ACM SIGCOMM Computer Communication Review* 55, 2 (2025), 13–21.
 - [31] Federal Communications Commission (FCC). 2022. Request for Orbital Deployment and Operating Authority for the SpaceX Gen2 NGSO Satellite System. Available at https://docs.fcc.gov/public/attachments/FCC-22-91A1_Rcd.pdf. Accessed on September 29th, 2025.
 - [32] Federal Communications Commission (FCC). 2024. Spectrum Sharing Rules for NGSO Fixed-Satellite Service Systems. Available at <https://www.federalregister.gov/documents/2024/12/13/2024-28993/spectrum-sharing-rules-for-ngso-fixed-satellite-service-systems>. Accessed on September 29th, 2025.
 - [33] Jose Flécher. 2024. Digicel leaves its customers in Haiti stranded as thousands left without internet. Available at <https://haitiantimes.com/2024/09/18/digicel-leaves-its-customers-stranded/>. Accessed on September 29th, 2025.
 - [34] SubTel Forum. 2013. BTC-selects-huawei-marine-to-upgrade-bahamas-domestic-submarine-network-international. Available at <https://subtelforum.com/75btc-selects-huawei-marine-to-upgrade-bahamas-domestic-submarine-network-international/>. Accessed on September 29th, 2025.
 - [35] Johan Garcia, Simon Sundberg, and Anna Brunstrom. 2025. A Detailed Characterization of Starlink One-way Delay. In *Proceedings of the 2025 3rd Workshop on LEO Networking and Communication*. 43–49.
 - [36] Jack Haddon. 2024. Ghana to licence Starlink in response to subsea cable cuts. (Accessed on September 29th, 2025).
 - [37] The New Zealand Herald. 2022. Elon Musk's SpaceX to help Tonga fix internet with Starlink satellite technology. Available at <https://www.scmp.com/news/asia/australasia/article/3166029/elon-musk-spacex-help-tonga-fix-internet-starlink-satellite>. Accessed on September 29th, 2025.
 - [38] Jeremy Hsu. 2025. NATO tests satellite internet as backup to sabotaged undersea cables. Available at <https://www.newscientist.com/article/2464396-nato-tests-satellite-internet-as-backup-to-sabotaged-undersea-cables/>. Accessed on September 29th, 2025.
 - [39] International Telecommunication Union (ITU). 1992. Constitution and Convention of the International Telecommunication Union.
 - [40] International Telecommunication Union (ITU). 2024. Radio Regulation: Articles.
 - [41] Liz Izhikevich, Manda Tran, Katherine Izhikevich, Gautam Akiwate, and Zakir Durumeric. 2024. Democratizing LEO Satellite Network Measurement. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 8, 1 (2024), 1–26.
 - [42] Jonathan McDowell. 2024. Satellite Constellations - Orbital distribution. Available at <https://www.planet4589.org/space/con/star/planes.html>. Accessed on September 29th, 2025.
 - [43] Michael Kan. 2025. Starlink Offers Cheapest Plan Yet to Entice Inactive Users. Available at <https://www.pcmag.com/news/starlink-offers-cheapest-plan-yet-to-entice-inactive-users>. Accessed on September 29th, 2025.
 - [44] Mohamed M Kassem and Nishanth Sastry. 2024. xeoverse: A Real-time Simulation Platform for Large LEO Satellite Mega-Constellations. *arXiv preprint arXiv:2406.11366* (2024).
 - [45] Kassing, Simon and Bhattacharjee, Debopam and Águas, André Baptista and Saethre, Jens Eirik and Singla, Ankit. 2020. Exploring the "Internet from space" with Hypatia. In *Proc. of ACM IMC '20*. 214–229.
 - [46] T.S. Kelso. 2025. Celestrak. Available at <https://celestrak.org/>. Accessed on September 29th, 2025.
 - [47] Simon Kemp. 2024. Digital 2024: Lithuania. Available at <https://datareportal.com/reports/digital-2024-lithuania>. Accessed on September 29th, 2025.
 - [48] Eunsun Kim, Ian P Roberts, Taekyun Lee, and Jeffrey G Andrews. 2025. Satellite Selection for In-Band Coexistence of Dense LEO Networks. *arXiv preprint arXiv:2503.15262* (2025).
 - [49] Kontur. 2023. Kontur Population: Global Population Density for 400m H3 Hexagons. Available at <https://data.humdata.org/dataset/kontur-population-dataset>. Accessed on September 29th, 2025.
 - [50] Zeqi Lai, Hewu Li, Yangtao Deng, Qian Wu, Jun Liu, Yuanjie Li, Jihao Li, Lixin Liu, Weisen Liu, and Jianping Wu. 2023. StarryNet: Empowering Researchers to Evaluate Futuristic Integrated Space and Terrestrial Networks. In *USENIX Symposium on Networked Systems Design and Implementation (NSDI)*. USENIX.
 - [51] Dominic Laniewski, Eric Lanter, Simon Beainn, Jan Dunker, Michael Dückers, and Nils Aschenbruck. 2024. Starlink on the Road: A First Look at Mobile Starlink Performance in Central Europe. In *2024 8th Network Traffic Measurement and Analysis Conference (TMA)*. IEEE, 1–8.
 - [52] Daniele Lepido and Donato Paolo Mancini. 2025. Italy Plans \$1.6 Billion SpaceX Telecom Security Deal - Bloomberg. Available at <https://www.bloomberg.com/news/articles/2025-01-05/italy-plans-1-5-billion-spacex-telecom-security-services-deal>. Accessed on September 29th, 2025.
 - [53] Kevin T Li, Christian A Hofmann, and Andreas Knopp. 2025. Simulation and Comparison of Vehicle Satellite Connectivity under a 3D Foliage Environment: Starlink and OneWeb. In *Proceedings of the 2025 3rd Workshop on LEO Networking and Communication*. 15–20.
 - [54] Yuanjie Li, Lixin Liu, Hewu Li, Wei Liu, Yimei Chen, Wei Zhao, Jianping Wu, Qian Wu, Jun Liu, and Zeqi Lai. 2024. Stable Hierarchical Routing for Operational LEO Networks. In *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking* (Washington D.C., DC, USA) (ACM MobiCom '24). Association for Computing Machinery, New York, NY, USA, 296–311. <https://doi.org/10.1145/3636534.3649362>
 - [55] Zhiyuan Lin, Zuyao Ni, Linling Kuang, Chunxiao Jiang, and Zhen Huang. 2022. Multi-Satellite Beam Hopping Based on Load Balancing and Interference Avoidance for NGSO Satellite Communication Systems. *IEEE Transactions on Communications* 71, 1 (2022), 282–295.
 - [56] Paul Lipscombe. 2022. TCC signs agreement with Kacific to expand connectivity in Tonga. Available at <https://www.datacenterdynamics.com/en/news/tcc-signs-agreement-with-kacific-to-expand-connectivity-in-tonga/>. Accessed on September 29th, 2025.

- [57] Lixin Liu, Yuanjie Li, Hewu Li, Jiabo Yang, Wei Liu, Jingyi Lan, Yufeng Wang, Jiarui Li, Jianping Wu, Qian Wu, Jun Liu, and Zeqi Lai. 2024. Democratizing Direct-to-Cell Low Earth Orbit Satellite Networks. In *21st USENIX Symposium on Networked Systems Design and Implementation (NSDI 24)*. USENIX Association, Santa Clara, CA, 791–808. <https://www.usenix.org/conference/nsdi24/presentation/liu-lixin>
- [58] Amanda Macias and Michael Sheetz. 2023. Pentagon awards SpaceX with Ukraine contract for Starlink satellite internet. Available at <https://www.cnn.com/2023/06/01/pentagon-awards-spacex-with-ukraine-contract-for-starlink-satellite-internet.html>. Accessed on September 29th, 2025.
- [59] Doug Madory. 2023. Dual Subsea Cable Cuts Disrupt African Internet. Available at <https://www.kentik.com/blog/dual-subsea-cable-cuts-disrupt-african-internet/>. Accessed on September 29th, 2025.
- [60] Marlene H. Dortch, Federal Communications Commission (FCC). 2025. ICFS File Nos. SAT-LOA-20200526-00055, SAT-AMD-20210818-00105, SATAMD-20221216-00175; SAT-MOD-20241011-00224 and SAT-AMD-20241017-00228; Call Signs: S3069 and S2992. Available at https://www.viasat.com/content/dam/us-site/corporate/documents/SpaceX_Response_Gen2.pdf. Accessed on September 29th, 2025.
- [61] Alexander Martin. 2025. UK to examine undersea cable vulnerability as Russian spy ship spotted in British waters. Available at <https://therecord.media/britain-undersea-cables-russian-spy-ship>. Accessed on September 29th, 2025.
- [62] Alan Mauldin. 2023. Is the Lifespan of a Submarine Cable Really 25 Years? Available at <https://blog.telegeography.com/2023-mythbusting-part-2>. Accessed on September 29th, 2025.
- [63] Maya Mehrara. 2025. Starlink Usage in Iran Skyrockets. Available at <https://www.newsweek.com/starlink-usage-iran-skyrockets-break-internet-law-rebel-restrictions-2011111>. Accessed on September 29th, 2025.
- [64] Praveen Menon and Tom Westbrook. 2022. An undersea cable fault could cut Tonga from the rest of the world for weeks. Available at <https://www.reuters.com/markets/funds/undersea-cable-fault-could-cut-off-tonga-rest-world-weeks-2022-01-18/>. Accessed on September 29th, 2025.
- [65] François Michel, Martino Trevisan, Danilo Giordano, and Olivier Bonaventure. 2022. A First Look at Starlink Performance. In *Proceedings of the 22nd ACM Internet Measurement Conference (Nice, France) (IMC '22)*. Association for Computing Machinery, New York, NY, USA, 130–136. <https://doi.org/10.1145/3517745.3561416>
- [66] Nitinder Mohan, Andrew E Ferguson, Hendrik Cech, Rohan Bose, Prakita Rayyan Renatin, Mahesh K Marina, and Jörg Ott. 2024. A Multifaceted Look at Starlink Performance. In *Proceedings of the ACM on Web Conference 2024*. 2723–2734.
- [67] Tamer Omar, Zakhia Abichar, Ahmed E Kamal, J Morris Chang, and Mohammad Abdullah Alnuem. 2016. Fault-Tolerant Small Cells Locations Planning in 4G/5G Heterogeneous Wireless Networks. *IEEE Transactions on Vehicular Technology* 66, 6 (2016), 5269–5283.
- [68] Nils Pachler, Edward F Crawley, and Bruce G Cameron. 2024. Flooding the Market: Comparing the Performance of Nine Broadband Megaconstellations. *IEEE Wireless Communications Letters* (2024).
- [69] Nils Pachler, Inigo del Portillo, Edward F Crawley, and Bruce G Cameron. 2021. An Updated Comparison of Four Low Earth Orbit Satellite Constellation Systems to Provide Global Broadband. In *2021 IEEE international conference on communications workshops (ICC workshops)*. IEEE, 1–7.
- [70] Nils Pachler de la Osa. 2024. *Optimizing resource allocation in large communications satellite constellations*. Ph.D. Dissertation. Massachusetts Institute of Technology.
- [71] Nils Pachler de la Osa, Markus Guerster, Inigo del Portillo Barrios, Edward Crawley, and Bruce Cameron. 2021. Static beam placement and frequency plan algorithms for LEO constellations. *International Journal of Satellite Communications and Networking* 39, 1 (2021), 65–77.
- [72] Sergei Pekhterev. 2021. The bandwidth of the StarLink Constellation... and the assessment of its potential subscriber base in the USA. *SatMagazine* 11 (2021), 54–57.
- [73] Winston Qiu. 2023. Amicitie Cable System Ready for Service. Available at <https://www.submarinenetworks.com/en/systems/trans-atlantic/amitie/amitie-cable-system-ready-for-service>. Accessed on September 29th, 2025.
- [74] Jason Rainbow. 2024. Starlink activates emergency phone connectivity for Hurricane Milton. Available at <https://spacenews.com/starlink-activates-emergency-phone-connectivity-for-hurricane-milton/>. Accessed on September 29th, 2025.
- [75] Chathurika Ranaweera, Mauricio GC Resende, Kenneth Reichmann, Patrick Iannone, Paul Henry, Byoung-Jo Kim, Pete Magill, Kostas N Oikonomou, Rakesh K Sinha, and Sheryl Woodward. 2013. Design and optimization of fiber optic small-cell backhaul based on an existing fiber-to-the-node residential access network. *IEEE Communications Magazine* 51, 9 (2013), 62–69.
- [76] Reuters. 2024. Israel eyes use of Musk's Starlink in event of war with Hezbollah, Calcalist reports. Accessed on September 29th, 2025.
- [77] Denys Rozenvasser and Kateryna Shulakova. 2023. Estimation of the Starlink Global Satellite System Capacity. In *Proceedings of the 12th International Conference on Applied Innovation in IT (ICAII 2024)*. 55–59.
- [78] Mario A. Sanchez, Fabian E. Bustamante, Balachander Krishnamurthy, Walter Willinger, Georgios Smaragdakis, and Jeffrey Erman. 2014. Inter-Domain Traffic Estimation for the Outsider. In *Proceedings of the 2014 Conference on Internet Measurement Conference (Vancouver, BC, Canada) (IMC '14)*. Association for Computing Machinery, New York, NY, USA, 1–14. <https://doi.org/10.1145/2663716.2663740>
- [79] Kanishka Singh. 2024. Dozens of Starlink systems deployed for Hurricane Helene by Biden administration. Available at <https://www.reuters.com/world/us/trump-says-he-spoke-musk-about-getting-internet-areas-hit-by-hurricane-2024-09-30/>. Accessed on September 29th, 2025.
- [80] Mackenzie Smith. 2022. Pacific digital infrastructure under spotlight after Tonga disaster. Available at <https://www.abc.net.au/pacific/programs/pacificbeat/pac-satellite-internet/13729920>. Accessed on September 29th, 2025.
- [81] Internet Society. 2024. 2024 West Africa Submarine Cable Outage Report. Available at <https://www.internetsociety.org/resources/doc/2024-west-africa-submarine-cable-outage-report/>. Accessed on September 29th, 2025.
- [82] SpaceX. 2016. FCC Selected Application for Space Exploration Holdings, LLC. Available at <https://fcc.report/IBFS/SAT-LOA-20161115-00118/1158350.pdf>. Accessed on September 29th, 2025.
- [83] SpaceX. 2019. SPACEX NON-GEOSTATIONARY SATELLITE SYSTEM. Available at <https://fcc.report/IBFS/SAT-MOD-20190830-00087/1877671>. Accessed on September 29th, 2025.
- [84] SpaceX FCC Filing. 2017. SpaceX V-BAND NON-GEOSTATIONARY SATELLITE SYSTEM. Available at <https://fcc.report/IBFS/SAT-LOA-20170301-00027/1190019.pdf>. Accessed on September 29th, 2025.
- [85] SpaceX Update. 2017. SPACEX NON-GEOSTATIONARY SATELLITE SYSTEM. Available at <https://fcc.report/IBFS/SAT-MOD-20181108-00083/1569860.pdf>. Accessed on September 29th, 2025.
- [86] Joe Supan. 2024. Starlink Puts Customers On Waitlist in Several US Cities. Available at <https://www.cnet.com/home/internet/starlink-puts-customers-on-waitlist-in-several-us-cities/>. Accessed on September 29th, 2025.
- [87] Hammas Bin Tanveer, Mike Puchol, Rachee Singh, Antonio Bianchi, and Rishab Nithyanand. 2023. Making Sense of Constellations: Methodologies for Understanding Starlink's Scheduling Algorithms. In *Companion of the 19th International Conference on Emerging Networking EXperiments and Technologies*. 37–43.
- [88] Karl Vick. 2023. Inside the Clandestine Efforts to Smuggle Starlink Internet Into Iran. Available at <https://time.com/6249365/iran-elon-musk-starlink-protests/>. Accessed on September 29th, 2025.
- [89] Vivek Wadhwa and Alex Salkever. 2022. How Elon Musk's Starlink Got Battle-Tested in Ukraine. Available at <https://foreignpolicy.com/2022/05/04/starlink-ukraine-elon-musk-satellite-internet-broadband-drones/>. Accessed on September 29th, 2025.
- [90] Caleb J. Wang, Ying Zhang, Qianli Dong, Esteban Carisimo, Ramakrishnan Durairajan, and Fabian E. Bustamante. 2025. Threading the Ocean: Mapping Digital Routes Across Submarine Cables using Calypso. In *Proceedings of the ACM SIGCOMM 2025 Conference (SIGCOMM '25) – Short Papers*. Association for Computing Machinery, Coimbra, Portugal. <https://doi.org/10.1145/3718958.3750512>
- [91] Shangguang Wang, Qing Li, Mengwei Xu, Xiao Ma, Ao Zhou, and Qibo Sun. 2021. Tiansuan Constellation: An Open Research Platform. In *2021 IEEE International Conference on Edge Computing (EDGE)*. IEEE, 94–101.
- [92] Shaowei Wang and Chen Ran. 2016. Rethinking cellular network planning and optimization. *IEEE Wireless Communications* 23, 2 (2016), 118–125.
- [93] Chris Woodhams. 2025. Undersea Vigilance: Britain monitors Russian spy ship amid heightened maritime security. Available at <https://boatingnz.co.nz/2025/01/undersea-vigilance-britain-monitors-russian-spy-ship-amid-heightened-maritime-security/>. Accessed on September 29th, 2025.
- [94] Niva Yadav. 2024. BW Digital to Expand Hawaiki Cable to Connect Tonga. Available at <https://subtelforum.com/bw-digital-to-expand-hawaiki-cable-to-connect-tonga/>. Accessed on September 29th, 2025.
- [95] Jingwei Zhang, Dali Qin, Chuili Kong, Feiran Zhao, Rong Li, Jun Wang, and Ye Wang. 2023. System-Level Evaluation of Beam Hopping in NR-Based LEO Satellite Communication System. In *2023 IEEE Wireless Communications and Networking Conference (WCNC)*. IEEE, 1–6.
- [96] Till Zimmermann, Eric Lanfer, Dominic Laniewski, Simon Brinkmann, and Nils Aschenbruck. 2025. Better Fill Up Your Pipe-Revisiting Starlink's Burst Characterization. In *Proceedings of the 2025 3rd Workshop on LEO Networking and Communication*. 36–42.

Algorithm 1 The GCB Terminal Distribution Algorithm

```

1: procedure GREEDYTD(cell_populations, num_terminals, cap=0)
2:   sorted_cells  $\leftarrow$  Sort cells by population density in descending order
3:   cell_terminals  $\leftarrow$  Empty dictionary
4:   terminals_left  $\leftarrow$  num_terminals
5:   for each cell in sorted_cells do
6:     if cell_population[cell]  $\geq$  cap then
7:       cell_terminals[cell]  $\leftarrow$   $\min(200, \text{terminals\_left})$ 
8:       terminals_left  $\leftarrow$  terminals_left - cell_terminals[cell]
9:     end if
10:  end for
     $\triangleright$  If any terminals are left, assign them uniformly amongst cells with more
    than pop_limit population
11:  return cell_terminals
12: end procedure

```

Algorithm 2 Greedy RF Beam Allocation

```

1: procedure GREEDYRF(cells, satellites, policy)
2:   beam_pool  $\leftarrow$  all beams (f, s, ch)
3:   map  $\leftarrow$  Empty dictionary
4:   sat_load[s]  $\leftarrow$  0 for all s
5:   cell_priorities  $\leftarrow$  values assigned by policy
6:   sorted_cells  $\leftarrow$  sort cells by cell_priorities
7:   while demand remains and beam_pool not empty do
8:     for each cell c in sorted_cells do
9:       cands  $\leftarrow$  beams from visible sats, ranked by capacity
10:      Pick feasible beam (f, s, ch) from cands that obeys interference rules
11:      if beam found then
12:        Assign (c, ch)  $\mapsto$  (f, s, ch); update loads and demand
13:        Prune interfering beams from beam_pool
14:      end if
15:    end for
16:  end while
17:  return map
18: end procedure

```

A Auxiliary Results

Results shown in figures 15 and 16 represent the capacity obtained in each of the case studies under wireless channel configurations Config B and Config C, respectively.

B Resource Allocation Algorithm

Algorithm 1 shows our uniform terminal distribution algorithm. The algorithm takes into account a cap on the minimum population size of cells that gets assigned terminals. Each cell is assigned 200 terminals in each iteration to maximize the maximize spectrum utilization in that cell.

Algorithm 2 shows our greedy RF beam allocation algorithm. The algorithm assigns beams to cells based on a priority ordering determined by either the spectrum allocation policy – either demand-driven or fairness-based. In each iteration, a candidate beam from a visible satellite is selected, and then assigned if it satisfies interference constraints. After each assignment, conflicting beams are pruned, and the process continues until either all demand is met or the beam pool is exhausted.

C Extensive Case Studies

This section presents the geographic distribution of our six case studies as shown in figure 17, and provides a detailed analysis of the relative importance of submarine cables for Tonga, Ghana, and South Africa. Additionally, it expands the discussion to three more case studies, offering a comprehensive evaluation. Figure 18 illustrates the relative importance of direct and indirect submarine

cables in Tonga, Ghana, and South Africa. To enhance visibility, submarine cables with a relative importance ratio of less than 1% are excluded from the ranking figure.

Great Britain - August 2024. Great Britain’s submarine cable system is a critical part of its national infrastructure, supporting both telecommunications and energy transmission [23]. Nearly 99% of Britain’s data traffic depends on 56 active undersea cables [10, 61]. As a key global connectivity hub, the country relies on this extensive network to facilitate data exchange between Europe, North America, and other regions [6]. However, this dependence on international links introduces significant risks, as disruptions could have severe economic and security consequences. Recent concerns over potential sabotage have led to increased surveillance of undersea infrastructure, particularly in the North Sea and the English Channel, amid rising geopolitical tensions [93].

On August 1, 2024, Telegeography reported damage to the AMITIE cable, which failed at segment 1.2 within the UK Exclusive Economic Zone (EEZ). However, no further details were provided regarding its impact. One possible reason for the limited disruption is Great Britain’s extensive submarine cable infrastructure. Analyzing one month of traceroute data, we observed 75% of the routes traversed submarine cables. As shown in Figure 19a, AMITIE ranks eighth among numerous submarine cables, indicating relatively low criticality in the network.

Despite this resilience, concerns are growing over the aging submarine cable infrastructure. According to Telegeography, 18 cables currently in service were activated before 2000, meaning they have been operational for over 25 years—well beyond the average submarine cable lifespan of 17 years [62]. Compared to other countries analyzed in case studies, even after excluding cables nearing decommissioning, the UK still has a substantial number of operational submarine cables. However, proactive measures are needed to support increasing capacity demands.

The impacted cable has fully commissioned since July 2023 and has an estimated total capacity 400Tbps [73].

Haiti - September 2024. Haiti, an island nation in the Caribbean, is highly vulnerable to natural disasters such as hurricanes and earthquakes due to its geographic location. The country relies heavily on a small number of submarine cables, including FIBRALINK and BAHAMAS DOMESTIC SUBMARINE NETWORK (BDSN), to maintain its domestic and international connectivity. With limited infrastructure, these cables are critical for supporting Haiti’s data and communication needs.

Since September 14, a damaged submarine cable near Kaliko Beach has led to significant internet connectivity problems for thousands of Digicel customers in Haiti. This disruption has severely affected families in the metropolitan area of Port-au-Prince, depriving them of reliable service and worsening the difficulties of living in a society of gang-related insecurity[33]. Digicel, Haiti’s largest provider of mobile and internet services, pinpointed the severed cable FIBRALINK, situated roughly 34 miles north of Port-au-Prince, as the primary cause of the outage.

Analyzing one month of traceroute data, we observed 82% of the routes rely on submarine cables. Figure 19b illustrates the relative importance of submarine cables in Haiti, highlighting the country’s heavy reliance on FIBRALINK for international traffic. With a design




Figure 15: Even with a different wireless model with reducing the Ku band data rate to 0.956 Gbps (config B in Table 2), the estimated capacity for various terminal distribution strategies follows similar trends.




Figure 16: Even with a different wireless model with increasing the Ku band data rate to 2.5 Gbps (config C in Table 2) and increasing the number of satellites to 14k, the estimated capacity for various terminal distribution strategies follows similar trends.




Figure 17: Geographic Distribution of Case Studies

capacity of only 320 Gbps [1], FIBRALINK poses a significant bottleneck, as Haiti's lack of redundancy makes even minor disruptions highly consequential. The only other significant submarine cable, BDSNI, now offers nearly twice the capacity of FIBRALINK after a recent upgrade [34]. However, limited overall capacity and aging

infrastructure continue to challenge Haiti's ability to meet growing connectivity demands, especially as digital services become increasingly vital for communication, commerce, and emergency response.

Lithuania – November 2024. Lithuania, a Baltic nation with a population of approximately 2.8 million, relies on a combination of terrestrial and submarine cables to sustain its international internet connectivity. The BCS EAST-WEST INTERLINK, a 530-kilometer (330-mile) submarine cable system, connects the Baltic Sea to the North Sea, serving as a critical link for Lithuania's internet traffic to Western Europe and beyond.

The disruption of the BCS EAST-WEST INTERLINK on November 17, 2024, resulted in a significant reduction of Lithuania's available bandwidth by approximately 30%, necessitating immediate rerouting efforts through alternative infrastructure [13]. While terrestrial routes provided temporary relief, the incident underscored the critical role of submarine cables in ensuring long-term connectivity




Figure 18: Submarine Cable Relative Importance in Tonga, Ghana, and South Africa




Figure 19: Submarine Cable Relative Importance in Great Britain, Haiti, and Lithuania

resilience. Our analysis processed a total of one month traceroutes, with 65% traversing the submarine cables. As shown in figure 19c, terrestrial paths are ranked as the most frequently utilized routes, based on observed traffic patterns and inferred mappings. However, in scenarios where submarine cables are employed, the BCS East-West Interlink consistently ranks as a critical pathway, highlighting its indispensable role in enhancing Lithuania's network resilience and global connectivity.

BCS EAST-WEST INTERLINK cable has been in service since 1997, but no official data is available on its total capacity. However, reports indicate that during its failure, the cable carried approximately one-third of Lithuania's internet traffic. The most recent per-user bandwidth data, reported in 2016, estimated Lithuania's per-user bandwidth at 125.45kbps [28]. Since no updated information is available, we use this as a reference to estimate the country's internet capacity in 2024. With 2.41 million internet users in Lithuania as of January 2024 [47], the estimated total national capacity is approximately 302Gbps. Based on this, the BCS EAST-WEST INTERLINK cable is estimated to have a capacity of around 101Gbps.

D Starlink Configuration

Table 3 shows shells for the existing and planned Starlink constellation that we used for our analysis.

Figure 20 shows the motif we apply to create a mesh from three ISLs in contrast with +Grid connectivity. +Grid stipulates that every satellite has an ISL with a neighbor from both the adjacent orbital shells. We modify this so that in every orbital plane, satellites alternate between having a link with only one neighbor from the right or left orbital plane. While the exact motif for three-ISLs is unknown, we choose this one as it closely resembles +Grid providing all its stability benefits and it tiles uniformly for the entire constellation.

Shell	Alt.	Incli.	Orbits	Satellites	ISLs	Status
S1	550	53°	72	1584	0	Deployed
S2	540	53.2°	72	1584	3	Deployed
S3	570	70°	36	720	3	Deployed
S4	560	97.6°	6	348	3	Deployed
S5*	530	43°	28	2128	3	Deployed
S5*	530	43°	28	1232	3	Approved
S6	525	43°	28	3360	3	Approved
S7	535	33°	28	3360	3	Approved
S8	340	53°	48	5280	3	Pending
S9	345	46°	48	5280	3	Pending
S10	350	38°	48	5280	3	Pending
S11	360	96.9°	30	3600	3	Pending
S12	604	148°	12	144	3	Pending
S13	614	115.7°	18	324	3	Pending

Table 3: Configuration of the current and planned Starlink constellation [31, 82–85]




Figure 20: Depicting connectivity in a satellite network with three ISLs. With the +Grid connectivity, both the red and blue inter-orbit ISLs are used while for three ISLs, only the blue ones are used.

E Performance degradation following cable cut

RIPE Atlas probe 1003709, located in South Africa, recorded a marked increase in latency, jitter, and packet loss during the outage, shown in Figure 21.




Figure 21: Impact on latency, jitter, and packet loss during the August submarine cable failure in South Africa as seen by a RIPE probe.




Figure 22: Estimated capacity comparing max flow and hot potato for greedy heuristic based terminal deployment.




Figure 23: Impact on deprioritized incumbent traffic during Ghana's failover using 10k threshold.

F Further Failover Capacity Analysis

Impact of Routing Policy. Figure 22 shows the normalized capacity provided by the two traffic engineering policies. Max flow provides significantly higher capacity in most cases since it utilizes multiple gateways, ensuring that all the requested capacity is met. In contrast, hot-potato tends to route to one or two gateways, causing congestion on the links to those gateways resulting in significantly lower capacity.

Localized Incumbent Impact of Ghana's failover traffic. Figure 23 shows the impact of Ghana's failover traffic on the nearby region. The four countries in close proximity of Ghana using Starlink, namely, Sierra Leone, Liberia, Benin, and Nigeria are affected due to Ghana's failover traffic. Benin and Nigeria have a larger




Figure 24: Estimated capacity for different incumbent demands per satellite.




Figure 25: Screenshot of the interactive map of the available capacity produced by the CosmoSim webpage.

impact due to GSL congestion caused by sharing more satellites with Ghana due to being more closer.

Interaction with incumbent traffic. Figure 24 shows the normalized failover capacities while competing with higher-priority incumbent traffic. There are three factors at play in this scenario: incumbent demand, failover demand, and LEO network access by neighboring countries. As incumbent demand grows, ISL capacity is exhausted, limiting the amount of capacity available to failover traffic. Countries with higher failover demand are more impacted by incumbent demand as they need more capacity. Finally, neighboring countries that access Starlink will not only congest ISL capacity but also GSL capacity, competing for satellite access with failover traffic. Thus, South Africa can achieve 3x higher failover capacity (absolute) than Lithuania at an incumbent demand of 10Gbps.

G CosmoSim Webpage

We integrate CosmoSim into an interactive webpage that allows users to choose their country of interest, their choice of policies,

and the number of terminals to be deployed. The webpage displays an interactive map that shows cells that received terminals within the country of interest. In addition, it shows available capacity to incumbents located in cells around the world. The webpage also supports showing a map of all gateway and their utilization. We show the screenshot of the first map in Figure 25.

H Gains from CosmoSim Optimizations

In this paper, we analyze nearly 1000 different scenarios each requiring a complete simulator run. In the naive case, this takes more

than 4-5 hours per data point. While our optimization to reduce the number of nodes brings this time down to less than an hour per data point, the reusability of the initial phases of CosmoSim for later phases, particularly the graph generation phase, further helps to reduce the total time for all the different simulations. For instance, the graph generation phase that can take up to 30 minutes is reused 60-80 times resulting in multiple hours of savings.

I Ethics

This work does not raise any ethical issues.