

A Characterization of Route Variability in LEO Satellite Networks

Vaibhav Bhosale, Ahmed Saeed, Ketan Bhardwaj, and Ada Gavrilovska

Georgia Institute of Technology

`vbhosale6@gatech.edu, {asaeeed, ketan.bhardwaj, ada}@cc.gatech.edu`

Abstract. LEO satellite networks possess highly dynamic topologies, with satellites moving at 27,000 km/hour to maintain their orbit. As satellites move, the characteristics of the satellite network routes change, triggering rerouting events. Frequent rerouting can cause poor performance for path-adaptive algorithms (e.g., congestion control). In this paper, we provide a thorough characterization of route variability in LEO satellite networks, focusing on route churn and RTT variability. We show that high route churn is common, with most paths used for less than half of their lifetime. With some paths used for just a few seconds. This churn is also unnecessary with rerouting leading to marginal gains in most cases (e.g., less than a 15% reduction in RTT). Moreover, we show that the high route churn is harmful to network utilization and congestion control performance. By examining RTT variability, we find that the smallest achievable RTT between two ground stations can increase by $2.5\times$ as satellites move in their orbits. We show that the magnitude of RTT variability depends on the location of the communicating ground stations, exhibiting a spatial structure. Finally, we show that adding more satellites, and providing more routes between stations, does not necessarily reduce route variability. Rather, constellation configuration (i.e., the number of orbits and their inclination) plays a more significant role. We hope that the findings of this study will help with designing more robust routing algorithms for LEO satellite networks.

1 Introduction

Low Earth Orbit (LEO) Satellite networks are emerging as an essential part of the future of global telecommunication, with pilot networks already in deployment and more planned [8, 10, 11, 15]. Satellites operating in a low Earth orbit provide low-latency communication, making them superior to terrestrial networks in some scenarios [38, 39]. At the same time, operating in low earth orbit inherently makes the satellite network very dynamic, with satellites orbiting the Earth every 100 minutes to maintain their orbits [3]. As a result, a satellite is only visible for a maximum of a few minutes to any single ground station. The highly dynamic nature of LEO satellite constellations introduces significant variability in the underlying network characteristics including the topology of the network, leading to frequent changes in the path characteristics.

In this paper, we provide a thorough characterization of route variability, highlighting the benefits, downsides, causes, and potential remedies. We characterize the variability in the lengths of paths (i.e., RTT variability) and the churn it creates in

the paths (i.e., rerouting frequency). Route variability can negatively impact route-adaptive algorithms (e.g., congestion control), and in turn the quality of experience of network users. In addition, high route churn complicates traffic engineering decisions. We posit that a deeper understanding of this variability can lead to better designs of the network layer and transport layer algorithms, improving the performance of satellite networks. We focus on networks equipped with inter-satellite links (ISLs), where data travel through satellite hops between a pair of ground stations, as they provide the lowest latency and constitute the future of LEO satellite networks [32, 85]. We look at three constellations, leveraging publicly-available information about satellite constellations, using the Hypatia simulator [46]. Concretely, our study looks into the following three aspects of route variability in LEO satellite networks.

First, we evaluate the pervasiveness of route churn and its impact (§4). We study paths selected by the shortest-path routing algorithm (i.e., paths that were deemed the shortest between two ground stations at some point during our simulation). We find that 15% of paths selected by the routing algorithm are used for less than 10 seconds. Further, we show that more than 50% of paths selected by the routing algorithm are used for less than half of their lifetime. This high churn can be considered as it affords the sender the lowest possible latency to the receiver, adopting the stance of “delay is not an option” [38]. However, we observe that such rerouting decisions lead to less than 15% RTT reduction in 50% of the cases, highlighting that *this high churn in routing may be unnecessary for many applications*. Further, we show that high variability can be harmful, causing significant deterioration in the performance of path-adaptive algorithms (e.g., congestion control algorithms). This negatively impacts the quality of experience for end users. Moreover, we find that greedily minimizing path length makes some paths hotspots with all traffic between a pair of cities flocking to them, underutilizing network capacity and creating potential security vulnerabilities [34].

Second, we observe that the smallest achievable RTT between a pair of ground stations can grow by up to $2.5\times$ as satellites move in their orbits. Further, we found that this RTT variability is predictable – it correlates with the location of ground stations, exhibiting a spatial structure (§5). In particular, we find the variability to be high only when the communicating ground stations are within 1500-3000km of each other and the travel direction between them isn’t along any of the orbital planes. To better understand this structure, we examine the building blocks of LEO satellite network paths: Inter-Satellite Links (ISLs) and Ground-Satellite Links (GSLs). We show that ISLs have stable properties that are not significantly affected by the motion of satellites, with intra-orbit ISLs having much larger lengths and inter-orbit ISLs. On the other hand, GSLs exhibit significant variability in their lifetime and length, making them the main source of variability. Through this analysis, we show that drastic changes in the building blocks of a route can significantly impact its length. For example, replacing an intra-orbit ISL with an inter-orbit ISL can significantly reduce the length of a route. Drastic changes are more likely if the travel direction between the communicating ground stations is not along any of the orbital planes. The impact of such changes is a function of the total path length (i.e., the longer the path, the smaller the impact of a change in its building blocks). Thus, we conclude

that this structure is determined by the relative position of the communicating pair of ground stations (i.e., the distance and angle of travel between them).

Finally, we show that RTT variability does not necessarily decrease by increasing the number of deployed satellites (i.e., by increasing path diversity) (§6). In particular, we measure RTT variability in a constellation as we add orbital shells. Simply adding orbital shells doesn't reduce RTT variability, with variability depending on the exact configuration of each shell (i.e., the number of satellites and orbits and the inclination of orbits). To evaluate the impact of constellation configuration, we compare two Starlink constellation configurations submitted to the FCC. We find that the more recent configuration introduces more RTT variability than the abandoned configuration.

2 Background

We start with a brief discussion of relevant background on LEO satellite networks needed to follow the rest of our study.

2.1 Overview of LEO Satellite Networks

A LEO satellite network or a constellation comprises thousands of satellites orbiting the Earth at an altitude in the range of 200-1600 km [24]. Due to their closeness to earth, LEO satellite networks can provide low-latency communication, potentially outperforming terrestrial networks [38, 39]. Several commercial LEO satellite networks have been announced, including SpaceX's StarLink, Amazon's Kuiper, and OneWeb. LEO satellite networks have already seen great demand with Starlink announcing they are already playing a key role as a reliable medium of communication (e.g., providing one of the most reliable means of communication in Ukraine [79] and Iran [31]). LEO satellite networks are also being considered for various purposes as part of the 3GPP's standardization effort for non-terrestrial networks (NTNs) [56]. Examples include StarLink's broadband service [10], OneWeb's mobile backhaul [9] and the agreement between Verizon and Kuiper to build a satellite-based backhaul for 5G networks [48]. Furthermore, LEO satellite networks are also opening up newer avenues such as providing WiFi connectivity in airlines and cruise ships [13, 58, 83]. The projected success of the current reincarnation of LEO satellite networks is driven by the reduced cost of building and launching such satellites [4, 22, 25, 59, 62].

In a LEO satellite network, satellites are placed in a number of shells, each consisting of a number of circular orbits or orbital planes at a constant altitude. Orbital planes 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




Fig. 1: An illustration of the inclination angles of the first, third and fourth shells of the Starlink constellation

	Altitude	Inclination	Orbits	Satellites
Starlink	550	53°	72	1584
	540	53.2°	72	1584
	570	70°	36	720
	560	97.6°	6	348
	560	53°	4	172
Kuiper	630	51.9°	34	1156
	610	42°	36	1296
	590	33°	28	784
Telesat	1015	98.8°	27	351
	1325	50.88°	40	1320

Table 1: Configurations of the three largest proposed constellations [50–52, 70–73, 77, 78]. Note that the configuration of StarLink’s network have changed compared to those mentioned in a recent study [46], making all their orbits at an altitude of less than 600km, lowering latency and increasing topology dynamics.




Fig. 2: An illustration of a LEO satellite network, showing the different types of links.

smaller inclination angles to provide greater coverage to densely populated areas [70]. Figure 1 shows an example with three different orbital planes, each at a different inclination. Similar orbital planes are equally spaced to form an orbital shell. Table 1 highlights these parameters for three of the largest proposed constellations.

2.2 LEO Satellite Networks Topology

Satellites communicate with each other through laser-based Inter-Satellite Links (ISLs) and communicate with ground stations using radio-based Ground-Satellite Links (GSLs) that operate in the Ku/Ka bands. 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. Figure 2 shows an illustration of a LEO satellite network.

The LEO satellite network topology is highly dynamic in nature owing to the rapid motion of the satellites. A LEO satellite travels at about 27,000 km/hr to maintain its orbit. Thus, a satellite is visible for a maximum 10-12 minutes from any point on earth. However, the satellite has to conform to the elevation angle bounds required for communication, limiting the accessibility time to a maximum 4.5

minutes. The exact amount of time that a satellite remains visible from a ground station depends on the altitude of the satellite. At a higher altitude, a satellite travels at slightly slower speeds, increasing the duration of its visibility.


Currently, satellites do not rely on ISLs and use ground station gateways [5, 6]. There have been some environmental concerns regarding laser-based ISLs, leading to the initial batch of Starlink satellites being launched without ISLs [40]. However, there have been recent launches of satellites with ISLs onboard, with more launches of similar satellites planned [32]. The current plan for LEO satellite networks is to carry traffic through ISLs to the ground station closest to the destination server [38, 47]. Relying on ISLs yields high data rates and low latency, provides better resistance to weather conditions and faces no regulations and a reduced risk of jamming [41, 66]. Therefore, in this paper, we focus on ISL-based networks, where traffic is routed through the satellite network from the source terminal closest to the sender to the destination terminal closest to the receiver.

The 2018 FCC filings by Starlink indicate the presence of 4 silicon-carbide communication components on every satellite [73], with recent work identifying them to be used as ISLs [20, 38, 46]. ISLs can be dynamically configured to connect satellites, allowing for the formation of many different topologies. The setup of an ISL can take between a few tens of seconds [69] to about a minute [84], during which the link cannot be used. The setup time for Starlink ISLs might be lower because the inter-satellite distances are smaller in the Starlink constellation. However, link setup won't be instantaneous, greatly reducing the utility of these links if reconfiguration was frequent. Hence, we assume static ISL configurations that require no ISL reconfiguration operations. In particular, we assume that satellites are connected following the so-called +Grid configuration where each satellite connects to two satellites in its own orbit, and with one satellite in each of the adjacent orbits. This configuration has been selected by the earlier work as the most likely configuration to be used in practice [20, 26, 37, 38, 53, 57, 67, 68, 80, 81].

The +Grid topology has many configurations, depending on how inter-orbit links are formed. The configuration of inter-orbit ISLs depends on the phase shift of orbital planes. The phase shift is a value between zero and one, determining the relative motion of satellites in adjacent orbits. At zero, all satellites with the same index in all orbital planes cross the equator at the same time. At one, satellite n in orbital plane p crosses the equator at the same time as satellite $n+1$ in orbital plane $p+1$. We use a phase shift of 0.5 as it very closely corresponds to the phase offset parameter used by Starlink [74] and can potentially provide good coverage of the Earth by uniformly distributing satellites in orbit [19]¹.

Now consider how inter-orbit ISLs should be formed in the presence of a phase shift between orbital planes. If a satellite is connected to its nearest neighbors in adjacent orbits in the presence of phase shifts, the resulting topology will be an inclined grid, providing poor east-west paths [38]. Alternatively, prior work [19, 38, 46] has intuitively argued for a slight variation of the +Grid configuration where a satellite

¹ Since the phase offset does not impact the stability of the GSL and ISL connections which we show to be the major reason for route variability in satellite networks, our results in this paper hold true for any chosen phase offset value.



(a) The CDF for the ratio between the path lengths achieved using our +Grid variation and the nearest-neighbor +Grid
 (b) The CDF for the ratio between the highest and lowest latencies observed for the two ISL variations.

Fig. 3: Characterizing the benefits of using the variation of +Grid over the nearest-neighbor +Grid ISL configuration

connects with a nearest neighbor in one adjacent orbit and with a phase shifted neighbor in the other orbit. This slight variation on the +Grid topology results in shorter east-west paths. Path 2 in Figure 16 is an example of paths created by that modified configuration. We verified this intuition with the following simulation study. We explain our simulation setup in §3.

We ran a simulation of the two +Grid variations for 100 minutes using the top 100 cities worldwide as source-destination pairs (total 4950 pairs). We measure the RTTs for all 4950 pairs every second to observe the variability inherent to the two choices. We first look at the ratio of the latencies observed for the two different variations at every second for all these source-destination pairs in Figure 3a. While the nearest-neighbor +Grid configuration has shorter paths in more than 40% of the scenarios, the maximum latency gain is just 43%. On the other hand, when the nearest-neighbor +Grid configuration has longer paths, its paths can be more than 5 times longer. To better understand the performance of the nearest-neighbor +Grid configuration, we look at the ratio of the maximum RTT and the minimum RTT between every pair of cities during the course of our simulation (Figure 3b). We observe that the nearest-neighbor +Grid configuration can have this ratio greater than 7 compared to about 2.7 for the variation to +Grid we use. Based on this simulation, we conclude that the nearest-neighbor +Grid configuration increases the magnitude and variance of the lengths of paths. Thus, for the rest of this paper, we use the +Grid variation employed in earlier work, leading to the configuration that minimizes path length variability.

2.3 Routing in LEO Satellites

Routing and traffic engineering in LEO satellite networks as a problem has seen a lot of interest over the past few decades. Earlier work has focused on providing performance guarantees, identifying paths with low latency while balancing the load between them. The typical approach is to adapt common traffic engineering techniques for the LEO satellite networks. For example, a recent proposal uses equal-cost multi-path (ECMP) techniques to distribute load between different paths, while also using obstacle-avoiding rectilinear Steiner trees (OARSTs) to ensure individual links are not overloaded [44]. Another proposal performs traffic splitting in a delay bound

manner wherein a path is randomly chosen from a set of paths within a certain delay bound of the shortest path while still favoring paths with lower latency [86]. Yet another load balancing-based scheme allows a congested link to provide signals to its neighbours to find alternate shortest paths that do not consist of the congested link [76]. Another work proposes the use of multi-protocol label switching (MPLS) to improve the quality of service in satellite networks by decoupling packet forwarding from the information carried in the IP header and periodically distributing routing information to the gateways [17]. Other recent work [38, 39, 46, 60] has also focused on using this idea of continuously computing the shortest path between destination pairs, and updating routing rules accordingly. Another strategy, towards the same goal, was to find the next hop moving towards the direction of the destination [42]. Satellite networks can also be viewed as Mobile Adhoc Networks (MANETs) with a more predictable structure. Variations of AODV [64] were proposed for satellite networks, leveraging location information to minimize delay and delay jitter [63]. All proposed algorithms covered in this brief survey rely on path length either primarily or partially in selecting routes. Thus, in this paper, we focus on shortest path algorithms which has also been the approach followed by similar studies [38, 39, 46].

3 Study Setup

Our study relies on simulations performed using the Hypatia framework [46] as the starting point, augmenting it with additional emulators as needed for the purposes of our study. We leverage Hypatia to generate the Two Line Element (TLE) information for satellites, a standard representation for satellite orbits containing the satellite identifier and orbit parameters [1]. Using that information, we are able to determine the ISL and GSL connectivity to conform with all the necessary physical requirements. Hypatia’s routing algorithm selects the shortest path between a pair of ground stations every 100ms. We use an interval of 1s to accelerate our simulations.² We use the Cesium [7] (a javascript library for visualizing 3D data and structures) interface provided by Hypatia to generate path visualizations. We do not use the packet-level simulator provided by Hypatia. Instead, we feed the delay measurements collected from Hypatia simulations into Mahimahi [61], a framework that enables running recorded traffic under emulated network conditions, to emulate route variability. Further, we use Pantheon [82], a platform developed for the evaluation and comparison of congestion control algorithms, on top of Mahimahi, allowing us to use the real implementation of all the studied transport layer algorithms. We use a modified version of Mahimahi available at [16] which allows us to vary the RTT values.

We predominantly use the first shell of the Starlink constellation, as it is the main fully-deployed constellation in practice. We also use the first shells from the Kuiper and Telesat constellations to show that our conclusions hold across constellations. We report our simulation results for a simulated period of 100 minutes for the Starlink and Kuiper constellations and 110 minutes for the Telesat constellation. Our choice of these intervals is because it allows us to capture the full orbit time of satellites in the

² It is telling that we are still able to show the impact of frequent routing even with a reduced frequency of route updates.

studied constellations, where their orbit times are 96 minutes for Starlink, 97 minutes for Kuiper, and 105 minutes for Telesat. We assume that the TLEs remain constant for the duration of our simulations which fall in the range of 100-110 minutes. This is a fairly reasonable assumption since the TLEs change every few days [2]. Most of our results use the global 100 most populous cities as the ground stations and consider all possible 4950 source-destination pairs. Prior work explicitly removed source-destination pairs that are close to each other, to highlight the low latency offered by satellite networks compared to terrestrial networks [20]. We include them since we are concerned with the broader context where satellite networks are used as the main communication infrastructure (e.g., in rural and disaster-affected areas). All key parameters used in our study such as the configuration of satellites, GSLs, the ISLs (discussed in §2) are based on publicly-available information in addition to using a popular state-of-the-art SGP4 model [43] to predict the position of satellites at every timestep of our simulation.

4 Route Churn is rife, unnecessary, and harmful

We start our study by assessing the extent and impact of route churn. We define a route as a sequence of satellite hops connecting a specific pair of ground stations. We define route churn as the change of routes (i.e., a change in one or more of the satellite hops forming the route). Our study considers routes that were picked by the shortest path routing algorithm (i.e., routes that were considered at some point the shortest between a pair of ground stations). Our goal is to identify the scenarios in which significant switching between paths occurs, the performance gains it leads to (if any), and its downsides. We use path and route interchangeably.

4.1 Route Churn is rife

We measure path churn using the following two metrics:

- *the lifetime of a path*: the duration a path remains valid (i.e., usable), allowing two specific communication ground stations to reach each other through the satellite hops of that path, and
- *the usage time of that path*: the duration a path is chosen by a routing algorithm to route traffic between the two communicating ground stations.

The lifetime of a path is determined by the topology dynamics. For instance, a satellite can remain in view of a ground station for just a few minutes, putting a cap on the lifetime of any single path. Thus, network dynamics naturally introduce churn. However, churn can be exacerbated by the routing algorithm. Consider that a routing algorithm ranks valid paths and picks the shortest path. As the topology changes, the ranks can change, leading the routing algorithm to potentially change its selected path, abandoning a valid path. Thus, a valid path might not be used for all of its lifetime. Therefore, the usage time of a path is determined by the decisions of the routing algorithm and the dynamics of the topology. Comparing the usage

time and the lifetime of paths helps us paint a picture of the amount of variability in the path selection.

We look at paths used by the all the source-destination pairs from the top 100 cities in the first shells of the Starlink, Kuiper and Telesat constellations. We focus on the two metrics defined above to identify any emerging patterns. Figure 4 shows the CDFs comparing the two values for the three constellations.

Looking at the lower tail of the distribution of the usage time, we observe that 15% and 20% of paths are used for less than 10 seconds in Starlink and Kuiper, respectively. This percentage is lower at around 8% in Telesat because it operates at a higher altitude, allowing satellites to remain in view for longer periods of time. Using paths for such a short duration can be detrimental to path-adaptive algorithms as we show later. The lifetime of paths has much higher values with the probability of the lifetime of a path being less than ten seconds is less than 5% for all constellations. Thus, it can be inferred that this volatility in path usage is actually caused by the volatility of the ranks assigned to paths. We also analyze the relationship between the usage time and lifetime of paths. In particular, we compute the ratio between the usage time and the lifetime for all studied paths (Figure 5). The figure shows that for the three constellations at least 50% of the paths are used for less than half of their lifetime.

4.2 Is route churn necessary?

The results shown in Figure 5 are for paths that are considered the shortest path between two ground stations at some point during their lifetime (i.e., their length/latency is not only acceptable but it's the best possible for some duration of time). This observation drove us to question the value of such high churn in paths. In other words, we try to answer the question: *is it necessary to switch between paths at a high frequency?*




Fig. 4: The CDF of the usage time and lifetime of paths for three different constellations






Fig. 5: The CDF for the ratio between the lifetime of a path and its usage time



(a) The CDF for the ratio between the shortest path length and the longest path length for a pair of ground stations, highlighting that for 70% of them abandoning a path yields a maximum of 25% reduction in latency.



(b) The CDF for the ratio between usage time and lifetime of the longest shortest path for a pair of ground stations, highlighting that 70% of them are abandoned for more than half of their lifetime.

Fig. 6: Characterizing the lifetime and benefits of abandoning longest shortest paths, showing that for the majority of cases they are abandoned for no significant gain.

High churn is justifiable if it yields significant performance improvements (i.e., a valid path is always abandoned for a much better path). If this churn leads to modest improvements in most cases, then high churn is perhaps unnecessary and routing algorithms should be designed to balance churn and performance. To assess the value of churn, we compute the maximum possible gain in performance that can be achieved by abandoning a path. Specifically, given a pair of ground stations, we evaluate the maximum achieved RTT and the minimum achieved RTT for the duration of our simulations. The minimum RTT reflects the length of the shortest path possible between two ground stations. The maximum RTT reflects the length of the *longest shortest path*.

The longest shortest path is the longest path selected by a routing algorithm to connect a pair of ground stations. Consider that as the topology changes, the composition (i.e., hops) and length of the shortest path between any two ground stations changes. The routing algorithm always selects the shortest possible path. Amongst all these paths, we focus on the longest one, calling it as the longest shortest path. The ratio between the minimum RTT and the length of the longest shortest path reflects the highest performance gain that a routing algorithm can make when abandoning a path (Figure 6a). The figure shows that the maximum performance gain is less than 25% for 70% of the source-destination pairs. Note that this result is fairly conservative since we focus on the best possible performance high churn can produce. In many cases, abandoning a path would lead to smaller gains.

We contrast the result in Figure 6a with the ratio between the lifetime and usage time of longest shortest paths. Figure 6b shows that 70% of longest shortest paths are used for less than half of their lifetime (while 70% of them are only 25% longer than the best possible RTT as shown in Figure 6a). This implies that even if longest shortest path were to be abandoned for the maximum possible gain, that gain in most cases will be modest. Achieving the lowest possible latency matters for some applications (e.g., High-Frequency Trading [65]). However, it won't impact the performance of most applications, especially given that the latency of ISL-based LEO satellite networks can be around 30% better than the latency of the terrestrial Internet [38, 39].




Fig. 8: An illustration of the path utilization experiment between 2000 ground stations in New York and 2000 ground stations in London. The illustration shows the available paths between the two cities.

To better contextualize the result, we consider a concrete example. In particular, we consider the Jakarta-Bogotá route. Figure 7 shows a time series of RTT values. The thick grey line shows the actual RTTs that will be observed using the shortest path routing policy, whereas the dotted lines represent the RTT of individual paths for the time they are valid. While the first switch takes place at 17 seconds due to the end of the first path, the second switch takes place 5 seconds later due to a difference of 0.005 ms in the latencies of the second and third paths, only to switch back to the second path 8 seconds later due to the end of the third path. Such frequent switching leads to four changes in the first 42 seconds, with the maximum latency gain of about 3.5 milliseconds (i.e., about 2.5% of the total RTT).




Fig. 7: The RTT of paths between Jakarta and Bogotá, showing eight path changes in 200 seconds. Dotted lines represent the RTT of different paths (in different colors). The solid line represents the achieved RTT.

Takeaway: Path length variability causes a high churn in routes, yet the variability is very small in most cases that it doesn't warrant the high churn.

4.3 Route churn can be harmful

Impact of Path Variability. One can argue that any improvement in latency is worth the trouble. We don't disagree. However, we show that having such a high churn rate for paths can be very harmful to performance. First, consider the impact of aggressively minimizing path length on network utilization. Our observation is that greedy

selection of shortest paths drives all nodes in the same locale to flock to the same set of paths. To illustrate a problem, we create a hypothetical scenario where 2000 nodes in New York are attempting to communicate with 2000 nodes in London. The nodes in each city are uniformly distributed over a circle of 22 km radius as shown in Figure 8. We observe the paths taken by each of the 2000 connections and find that all connections flock to the same path, despite having other valid paths with marginally worse latency. Figure 9 shows the behavior. The thickness of the solid lines represents the number of connections with a certain RTT value. It's clear in this example that all connections are using the same paths. The purple path is abandoned by all connections for the yellow path for about half a millisecond lower RTT. All the connections abandon the yellow path for the black path for a half millisecond gain, only for all of them to reuse the yellow path after 28 seconds. These gains are considerably low compared to the RTT of the path (around 1-2%). In addition, the benefits of these gains may get diluted due to the inability of the transport layer to keep up with the changes.




Fig. 9: The RTT of paths between New York and London. Dotted lines represent the RTT of different paths (in different colors). The width of the solid line represents the achieved the number of ground station pairs taking a path. All pairs take the same path.

Next, we consider the impact of path churn on the behavior of the transport layer. We consider the link utilization, the 95th percentile delay, and the power defined as the ratio of utilization to the 95th percentile delay exhibited by different congestion control protocols. We look at the route between Pune, India and Lahore, Pakistan. A path change in a LEO satellite network may end up changing the observed RTT and bandwidth (e.g., switching to a path with a different number of flows competing for its bandwidth). Thus, we evaluate the impact of RTT variability and bandwidth variability, separately and combined. Figure 10 shows the delays for the 60 seconds time interval we use for this experiment. Instead of assigning a different bandwidth value for every path, we select two bandwidth levels that we alternate between with every path change to show the impact of changing bandwidth on TCP algorithms. In particular, the bandwidth changes between 204 Mbps and 48 Mbps. We choose these values as they closely correspond to the range of bandwidth specifications for starlink [75]. The results are shown in Table 2.




Fig. 10: The RTT of the route Pune and Lahore

As expected, there is a clear tradeoff between latency and utilization, with none of the studied algorithms being able to reach a good balance between the two. For example, BBR [23] achieves high bandwidth utilization, while introducing significant

		BBR	Cubic	PCC-Allegro	PCC-Vivace	Vegas
Constant	Utilization	95.4	61.8	93.5	86.8	10.1
	Delay	18.3	16.7	16.3	16.2	13.6
	Power	5.21	3.7	5.74	5.36	0.74
Variable	Utilization	90.1	55.4	70.6	64.6	45
	Delay	29.7	28.7	26.1	19.6	19.5
	Power	3.03	1.93	2.7	3.3	2.3
Constant RTT	Utilization	95.9	68.6	71.3	73.8	37.8
	Delay	26.4	25.1	16.5	16.0	16.3
	Power	3.63	2.73	4.3	4.61	2.32

Table 2: The utilization (%), 95th percentile one-sided delay (in ms from sender to receiver), and power (defined as the ratio of utilization to the 95th percentile delay) for three different scenarios of path variability for five different congestion control algorithms. **Green** reflects the best result in its row and **red** reflects the worst result in the row. No single algorithm optimizes both delay and utilization.

delay. On the other hand, Vegas [21] consistently achieves the lowest utilization with fairly low delays. PCC-Allegro [27] and PCC-Vivace [28] generally show utilization in the range 60–85%, while incurring fairly lower delays. This variation is further exacerbated when we look at the power values since both BBR and Vegas belong to the lower percentile of the spectrum with PCC-Allegro and PCC-Vivace outperforming on this metric. We recognize that earlier work in congestion control attempts to effectively handle bandwidth variability in datacenter networks [36, 55] and wireless last miles [35]. However, these solutions are typically tailored for their target network, motivating such tailoring for LEO satellite networks.

Takeaway: High churn in routes, caused by shortest path algorithms, can cause poor path utilization and poor performance by congestion control algorithms.

5 Understanding RTT Variability

High churn rates for paths are caused by frequently abandoning paths to reduce route latency. We have shown that for 70% of the studied ground station pairs the gains are smaller than 25% (Figure 6a). Route churn can be significantly reduced if the routing algorithm was made less aggressive (e.g., only abandoning a path if significant RTT gains are made). However, high churn can still happen when RTT variability is high (e.g., it is possible to improve the RTT by more than 25% if a path is abandoned). In this section, we identify the causes of high RTT variability. Our goal is to provide insights that can aid the design of better routing algorithms that can reduce route churn in the presence of RTT variability.

5.1 RTT Variability Exhibits Spatial Structure

As shown in Figure 6a, the smallest achievable latency for the same route can exhibit a high variability potentially introducing over $2.5\times$ higher RTT. We sought




Fig. 11: Heat maps showing the ratio between max RTT and min RTT in paths between Null Island(0° latitude, 0° longitude), Darfur, and Kyiv and 2700 nodes uniformly distributed around the globe. The redder the point, the higher the ratio, indicating higher variability.

to identify any patterns in this variability and observed the presence of a spatial structure correlating to the location of the ground stations. We measured RTT variability using the ratio between the maximum RTT and the minimum RTT observed when a shortest-path routing algorithm is used. Our simulation considers RTT values observed during a period of 100 minutes, using the first shell of the StarLink network. We report RTT variability in paths between 2700 uniformly distributed hypothetical destination ground stations and three source ground stations: Null Island (0° latitude, 0° longitude), Kyiv ($50.4501^\circ N, 30.5234^\circ E$), Ukraine, and Darfur ($14.3783^\circ N, 24.9042^\circ E$), Sudan, each capturing a different latitude. Figure 11 shows the results, where the darker colors represent higher variability. The darkest represents cases where the max RTT is more than $2\times$ the min RTT.

Figure 11 shows that there is a clear structure for ground station placements that would yield high variability. For example, low latitude source ground stations (e.g., Null Island and Darfur) observe high variability when communicating with ground




Fig. 12: The CDFs of different properties of ISLs

stations places in a ring-like structure with diagonal ribbons extending from it. On the other hand, the structure only includes the ribbons for high latitude stations (e.g., Kyiv). This structure only impacts destinations that are within 1500-3000 km from the source (geodesic distance). We found these results to hold regardless of the longitude of the source station and similar structures repeat for source stations at the same latitude.

We find that a particular route between two ground stations shows high variability when the makeup of the paths changes drastically as satellites move. To better understand this behavior, we study the building blocks for paths and their characteristics.

5.2 Building blocks for Paths : ISLs and GSLs

Every path from a terrestrial source to a terrestrial destination is comprised of two Ground-Satellite links (GSLs) and zero or more Inter-Satellite links (ISLs). For the +Grid topology we consider in this paper, there are two main categories of ISLs that differ significantly in their properties: inter-orbit ISLs and intra-orbit ISLs. The lifetime and length of a path are governed by the properties of its links. Thus, we take a closer look at the properties of different types of links. In particular, we examine the *length* and *lifetime* of those components.

Properties of ISLs. We record the lengths of all ISLs in a 100 minute time interval for the first shell of Starlink. We observe that the lengths of ISLs are highly predictable. Figure 12a shows the CDF of the median length of ISLs, broken down based on its type.

The results show that there are two types of inter-orbit ISLs: one with a median length of about 760 km and the other at 1384 km. The two types of inter-orbit ISLs occur alternately such that each satellite has one inter-orbit ISL of both the types. This is an outcome of the phased orbit structure we discussed in Section 2. On the other hand, intra-orbit ISLs are uniform with all their lengths at about 1970 km, which is about 150% more than the first cluster of the inter-orbit ISLs and 50% more than the second. There is little variability in the lengths of ISLs. Figure 12b shows the CDF of the ratio between the minimum length and maximum length of an ISL during the 100 minute period for the two types of ISLs. All intra-orbit ISLs and 50% of inter-orbit ISLs exhibit minimal length change (0.2% and 6%, respectively).

The length of the other 50% of inter-orbit ISLs can change by up to 21% due to the varying distances in between different orbital planes across latitudes [30, 45]. The orbits are closer to each other at higher latitudes compared to the lower ones, and hence the inter-orbit ISL lengths vary accordingly.

Takeaway: The length of a single ISL is stable and there are three different types of ISLs each having significantly different length.

Properties of GSLs. To study the variability in GSLs, we recorded the lengths of all GSLs from a ground station placed in Tokyo as a representative example for the same 100 minute interval. GSLs exhibit considerably different characteristics compared to the ISLs. They exhibit great variability in their lengths and lifetimes. Figure 13 shows the lifetime of different GSLs that a ground station in Tokyo can form in the 100 minute period. The figure shows that the lifetime of a GSL can be as low as 6 seconds, with a maximum of 4.5 minutes.

Figure 14 shows the lengths of all the GSL lengths in the 100 minute interval for Tokyo. The CDF captures the change in the length of individual GSLs as well as the variability between different GSLs. The lengths of GSLs are uniformly distributed between a minimum of 550 km to a maximum of 1254 km, depending on where the satellite lies with respect to the ground station. A satellite right above the ground station will have the shortest distance whereas a satellite at the periphery of the region the ground station can communicate with will have larger distance, and as the satellites move inside this region, their length fluctuates.




Fig. 13: The CDF of GSL lifetimes




Fig. 14: The CDF of GSL lengths.

Takeaway: Due to the stability of the ISLs, the variability in the lifetime of GSLs has a greater impact on the variability in the lifetime of paths.

Path formation Having looked at both the components individually, we now look at how they play a part in determining the length of paths. Intra-orbit ISLs travel along the orbital planes, whereas the inter-orbit ISLs travel between them. Thus, the makeup of a path depends on the direction of travel between the two communicating ground stations.




Fig. 15: An example showing that using an intra-orbit ISL when traversing across orbital planes can increase the path length

Communicating across orbital planes. We consider travel directions that crosses orbital planes. In such scenarios, the path can be made entirely of inter-orbit ISLs. Recall that the length of intra-orbit ISLs can be $2.7\times$ the length of inter-orbit ISLs. Thus, if a topology change forces replacing a single inter-orbit ISL with an intra-orbit ISL, the overall path length can change significantly. However, the impact of that change is also a function of the total length of the path. For example, replacing a single link in a path made of 20 links will result in much lower change in total path length compared to replacing a link in a path made of two links. To illustrate this point, consider paths from Kyiv to Cairo. Figure 15 shows two paths between the two cities. Each is the shortest available path during different time intervals. Figure 15a shows a path made entirely of three inter-orbit links. Figure 15b shows a path with a smaller hop count but with a single intra-orbit link, leading to a longer path due to the larger length of intra-orbit ISLs.

Takeaway: The length of paths communicating over short distances across orbital planes can be significantly increased by the addition of a single intra-orbit link.

Communicating along orbital planes. We consider an example where the direction of travel is along an orbital plane. In such scenarios, paths formed mostly by intra-orbit links yield the shortest paths. Despite having the largest lengths of all ISL types, intra-orbit ISLs help cover the distance along orbital planes along a straight line, providing shorter paths. Figure 16 shows an example for communication between Miami and Denver. Path 1 is the shortest path and it is formed exclusively by intra-orbit ISLs. We contrast with Path 2, where we picked a first hop with an orbital plane that is not parallel to the direction between the two cities. Despite being made up mostly of short inter-orbit ISLs, Path 2 is much longer than Path 1. It's




Fig. 16: An example showing the value of Intra-orbit links and the downside of not using them when applicable. Note that Path 2 is valid path but not a shortest path and is used just for illustration (i.e., never picked by a routing algorithm).




Fig. 17: The impact on the length of a path by using the worst first hop compared to the best one. The line shows the average of the ratio and the shaded part shows the standard deviation.

important to note that the GSL hops of Path 2 are both shorter and more long living than those of Path 1.

Takeaway: The length of paths communicating over relatively long distances along orbital paths can be significantly increased if inter-orbit communication is required.

Choosing GSLs. The choice of a GSL is not entirely a routing decision. Recall that GSLs are wireless links with SNR determining the quality of the link. The SNR depends on the length of the GSL and weather conditions among other things. Thus, the choice of GSLs can be made independent of the routing decision. We explore the impact of that choice on RTT variability. In particular, we assess the impact of the choice of GSLs, considering the worst case by examining first hops that yield the longest paths. In particular, we compute the ratio between the length of the shortest possible path through the worst case first hop, and the length of the actual shortest path. This metric measures the worst possible performance based on the choice of the first hop. We measure that metric for four source ground stations on the 85° longitude, uniformly spaced between a latitude of 5° and 55° . Each source ground

station communicates with 2700 points distributed uniformly within the coverage area of starlink’s first shell. Figure 17 summarizes the results as a function of the distance between communicating ground stations. The results show that a poorly selected first hop may double the RTT. However, with the increase in the distance between the source and the destination, this impact reduces.

Takeaway: The choice of the first hop is integral to determining and optimizing overall path length.

5.3 Darfur as a Case Study : Building Blocks in Action

As stated earlier, a particular route between two ground stations shows high variability when the makeup of the paths changes drastically. This can include changing the first hop orbit to be closer or farther away from the destination. It can also include changing the types of ISLs. We observe that these two phenomenon may occur in tandem which can further increase the length of the path by up to 2600 km, which could potentially double the length of the path. This also implies that the spatial structure is highly localized since the impact of changes in the makeup of a path decreases as the number of components of the path increases (i.e., as the distance between the communicating ground stations increases). Similarly, for points that are very close to each other, the path typically doesn’t include any ISLs, limiting variability.

We illustrate this using two exemplar routes originating in Darfur. We choose two destinations: Isangi($0^\circ N, 24^\circ E$), Democratic Republic of Congo that lies in the ring, and Muynak($44^\circ N, 60^\circ E$), Uzbekistan that lies outside the ring. We look at a 200 second time interval to observe the variation in lengths and the number of times different paths are chosen (Figure 18). The markers on each of the lines show when a path change occurred.




Fig. 18: The RTT variations for the two paths Darfur-Isangi and Darfur-Muynak with the markers representing path changes

The Darfur-Isangi route is a little awkwardly placed with respect to the topology of the Starlink’s first shell. The two ground stations are about 1600 km away which is low enough for them to be served by a single satellite. However, as the route doesn’t lie along any orbital plane, the distance isn’t low enough to always be served by a single satellite. Therefore, inter-orbit ISLs will be required for this path at certain instances, leading to high variations in RTT as we discussed earlier. We show three different paths at time steps 80, 81, and 86 in Figure 19. At time step 86, a single satellite is capable of communicating with both Darfur and Isangi and hence no other hops are needed. However, that is not the case at time step 80 and 81. At time step 80, the first hop uses an inter-orbit ISL to reach the next orbit and then an intra-orbit




Fig. 19: An example of extreme RTT variation showing three paths between Darfur and Isangi, reflecting time steps 80,81, and 86 in [Figure 18](#)




Fig. 20: An example of modest RTT variation showing two paths between Darfur and Muynak, reflecting time steps 42 and 43 in [Figure 18](#)

link to finally reach the destination. Whereas at time step 80, the first hop can reach the destination with just one intra-orbit ISL making it shorter.

Since the Darfur-Muynak route is along the orbital planes for the first shell, the intra-orbit ISLs are used predominantly. We highlight the switch happening at timestep 43 in [Figure 20](#). In this case, even swapping out an intra-orbit ISL for an inter-orbit one doesn't lead to much change since the links are always travelling towards the destination.




Fig. 21: The CDF of the maximum RTT variation observed by paths comparing the older and newer Starlink configurations

Altitude	Inclination	Orbits	Satellites
550	53°	72	1584
1110	53.2°	32	1600
1130	74°	8	400
1275	81°	5	375
1325	70°	6	450

Table 3: The shell parameters configuration for Starlink V1

6 Does Deploying More Satellites Reduce Variability?

We culminate our study by discussing the impact of increasing number of satellites on variability. LEO satellite network operators improve their presence and quality of coverage by launching more orbital shells. We compare the RTT variability for the first shell of Starlink with the entire Starlink constellation composed of five shells. We use both the Starlink constellation configurations released to the FCC. For brevity, we refer to the first shell as Starlink S1, and the older and newer proposals as Starlink V1 and Starlink V2 respectively. We simply analyze the impact of these configurations without claiming to have any information about the reasons for the change.

Since each satellite currently has only 4 lasers to support 4 ISLs, we assume that they are utilized to setup a +Grid leading to individual shells operating independently. ISLs only connect satellites belonging to the same shell. Different shells operate at different altitudes and their configurations include the number of orbits, the inclination of orbits, and the number of satellites per orbit. In these simulations, we use the global 100 most populous cities as the ground stations and look at all possible 4950 source-destination pairs. We record the RTT variations for both the Starlink configurations.

We consider the performance of the proposed Starlink configurations in Figure 21. We consider the older proposal as it has four different inclination values, while the current plan has only three, creating a missed opportunity in making use of inclination diversity. Further, it demonstrates the impact of varying altitudes significantly between shells in addition to varying inclination angles. The earlier plan had the last four shells at considerably higher altitudes compared to the current plan. Deploying satellites at higher altitudes makes them visible from ground stations for longer periods of time, reducing the pace of change in the topology. Thus, higher altitude shells reduce RTT variability. This can also be observed in the results presented in Section 4, comparing the first shell of Starlink and the first shell of Telesat.

Thus, we can infer that simply increasing the number of satellites doesn't help solve this problem. With approximately the same number of satellites, Starlink V1 can provide improved performance over Starlink V2 due to higher altitudes and a greater diversity in the inclination angles.

In addition to comparing the two publicly available configurations by Starlink, we also tried to individually profile the number of shells and the impact of diversity




Fig. 22: The CDF of the maximum RTT variation observed by paths for different shell configurations

in inclination angles. First, in Figure 22a, we study the impact on RTT variation by simply increasing the number of shells with similar orbital characteristics - inclination angle, number of orbits, and number of satellites as the first shell of Starlink. We observe that adding an extra shell improves the RTT variability a little, but after that, there is no tangible difference after adding two or six more shells leading us to conclude that simply replicating the shells will provide minimal to no gains.

Second, we also considered the impact of changing the inclination angles between shells. We repeat the earlier simulation while tweaking the inclination angles. We use 53° , 27° , 72° , 13° , 40° , 62° , 82° , 53° as the inclination angles in order³. We observe in Figure 22b that increasing the number of shells while varying the inclination angles can reduce RTT variability by $3\times$ in the median. To understand the role of varying orbit inclination, consider the Darfur-Isangi (Figure 19). Such source-destination pairs cause high RTT variability because the direction from the source to the destination does not lie on any orbital plane. Therefore, adding shells with different inclination angles reduces such occurrences, reducing RTT variability. Thus, our conclusion is that changing the inclination angle between shells yields significant benefits in reducing RTT variability.

7 Discussion

Topology variants. As discussed earlier, our simulations use a specific variant of the +Grid topology due to its good coverage properties, making it the most likely topology to be used in practice. However, there are several other inter-satellite network topologies with desirable properties. For example, extending inter-orbit ISLs beyond adjacent orbits was shown to reduce latency and improve network throughput [20]. This approach requires more ISL reconfigurations as satellites move. In our study, the properties of ISLs were relatively stable. Thus, we posit that any such added variability in ISLs can increase path churn by adding another source of variability. Future work should explore the tradeoffs offered by such topologies, taking into account variability, latency, and throughput.

Even for the +Grid topology, we focused on a single variant. The configuration of other variants will depend on the phase shift in orbits and how inter-orbit links are

³ Not all inclinations we used might be possible due to interference or orbital constraints. Our goal is simply to highlight the impact of that parameter.

formed. For example, a variant can create more uniform inter-orbit ISLs by connecting satellites to their nearest neighbors in adjacent orbits, harming latency but reducing path variability. We leave it for future work to explore the impact of such variants on route variability.

Routing Algorithm Variants. Path length is an important factor in all existing routing algorithms for LEO satellite networks. This motivated our study of path length variability and its impact on decisions made on shortest-path-routing. However, our study didn’t look into the impact of path length variability on routing algorithms that use more complicated metrics. We leave such studies to future work which we hope will come as a part of new routing metric proposals that mitigate the downsides of the variability identified in this study.

Ground relays. Our study does not take into account satellite networks that rely on ground relays. This is driven by the fact that currently deployed Starlink satellites use ground stations only to connect to the Internet directly and not to connect to each other. A network with ground relays, where ground stations and user terminals act as ground relays as described in [39], will likely exhibit a higher degree of variability than the network we studied owing to the increased number of GSLs which are a big contributor to route churn.

8 Related Work

Optimizing delay. The domain of LEO satellite networks has seen an increased amount of interest from the networking community over the past few years. The community has been especially excited about the potential of these networks to outperform terrestrial networks. This has led to topology design proposals that aim for inter-satellite network topology providing low latency [20, 38]. While these focused on ISL-based networks, others explored achieving low latency in the absence of ISLs using ground relays [39]. The goal for all of these studies is to optimize the network for low latency to outperform terrestrial fiber networks. We offer a different perspective, showing that optimizing exclusively for delay can be harmful to network utilization and path-adaptive algorithms. We also show that a slight sacrifice in delay can improve route stability. We hope that our insights will help design algorithms that can better navigate these tradeoffs. In addition, there have been multiple efforts [18, 33, 49] analyzing the challenges with integrating the LEO satellite network with the current internet backbone. They look at how satellites can be used to assist with inter-domain routing. While our work focuses only on intra-domain routing for the satellite network, it will help the design of such systems by providing better routing through the satellites.

Routing. A lot of work has been done in the past looking at different goals for routing such as reducing propagation delay [30, 38, 39, 46, 60], improved load balancing [76], and energy efficiency [14, 45] (§2). However, most of these proposals were made for an older generation of satellites and applications. The current generation consists of a considerably larger number of satellites and also incorporates many advancements in the satellite communications domain [12]. We hope our work motivates a resurgence in research on routing algorithms in satellite networks..

Variability. Earlier work tackles the impact of variability on addressing and configuration of routing tables [54]. It proposes a rethinking of the logical network topology to better accommodate the rapid changes in the physical infrastructure. While it can reduce the variability in the logical network, it doesn't consider the causes or remedies of variability in the physical network. Our work focuses on the physical network.

9 Conclusion

In this paper, we study the variability in paths rampant in LEO satellite networks. We concretely present the amount of route churn and RTT variability, also highlighting the impact of such variability on path utilization and congestion control. We delve deeper into the reasons why this variability exists by presenting the building blocks of paths and infer that this variability exhibits a spatial structure. Our hope is that this work will provide the key insights for the design of specialized routing and perhaps, congestion control algorithms for LEO satellite networks, taking into account that when delay is an option significant gains can be made in overall network performance.

Acknowledgments

This project was partially supported by NSF grant 2212098 as well as grants and gifts from Google, Cisco, the Georgia Smart program, and VMware. We thank CloudLab [29] for providing a testbed to run our simulations.

References

1. SGP4 Propagator. Available at https://help.agi.com/stk/11.0.1/Content/stk/vehSat_orbitProp_msgp4.htm
2. SGP4 Propagator. Available at https://help.agi.com/stk/11.0.1/Content/stk/vehSat_orbitProp_msgp4.htm (2016)
3. Popular Orbits 101. Available at <https://aerospace.csis.org/aerospace101/popular-orbits-101/> (2017/11/30) (2017)
4. Launch Costs to Low Earth Orbit, 1980-2100. Available at <https://www.futuretimeline.net/data-trends/6.htm> (2022/05/12) (2018)
5. FCC Selected Application for Space Exploration Holdings, LLC. SESLIC2019090601171 (2020)
6. FCC Selected Application for Space Exploration Holdings, LLC. SESLIC2019021100151 (2020)
7. Cesium: The Platform for 3D Geospatial. Available at <https://cesium.com/> (2023)
8. OneWeb. Available at <https://oneweb.net/> (2023)
9. OneWeb Mobile Backhaul. <https://oneweb.net/solutions/carrier-enterprise/mobile-backhaul> (2023)
10. Starlink. <https://www.starlink.com/> (2023)
11. Telesat. Available at <https://www.telesat.com/> (2023)
12. Ackerman, R.K.: Technology Drives New Satellite Communications Capabilities. Available at <https://www.afcea.org/content/technology-drives-new-satellite-communications-capabilities> (2020/08/01) (2020)

13. Airlines, H.: Hawaiian Airlines to Offer Free, High-Speed Starlink Internet Connectivity on Transpacific Fleet. Available at <https://newsroom.hawaiianairlines.com/releases/hawaiian-airlines-to-offer-free-high-speed-starlink-internet-connectivity-on-transpacific-fleet> (2022/10/28) (2022)
14. Akturan, R., Vogel, W.J.: Path diversity for leo satellite-pcs in the urban environment. *IEEE Transactions on Antennas and Propagation* **45**(7), 1107–1116 (1997)
15. Alan Boyle: Amazon to offer broadband access from orbit with 3,326 satellite 'Project Kuiper' constellation. Available at <https://www.geekwire.com/2019/amazon-project-kuiper-broadband-satellite/> (2019/04/04) (2019)
16. Ameya Daigavane: ameya98/mahimahi. Available at <https://github.com/ameya98/mahimahi/> (2019)
17. Berlioli, M., Donner, A., Menichelli, R., Werner, M.: Mpls traffic engineering for leo satellite constellation networks. In: 21st International Communications Satellite Systems Conference and Exhibit. p. 2327 (2003)
18. Bhattacharjee, D., Aqeel, W., Bozkurt, I.N., Aguirre, A., Chandrasekaran, B., Godfrey, P.B., Laughlin, G., Maggs, B., Singla, A.: Gearing up for the 21st century space race. In: *Proc. of HotNets '18*. pp. 113–119 (2018)
19. Bhattacharjee, Debopam: Towards Performant Networking from Low-Earth Orbit. Ph.D. thesis, ETH Zurich (2021)
20. Bhattacharjee, Debopam and Singla, Ankit: Network topology design at 27,000 km/hour. In: *Proc. of CoNext '19*. pp. 341–354 (2019)
21. Brakmo, L. S. and O'Malley, S. W. and Peterson, L. L.: TCP Vegas: New Techniques for Congestion Detection and Avoidance. In: *Proc. of ACM SIGCOMM '94*. pp. 24–35 (1994)
22. Brian Wang: SpaceX Starlink Satellites Could Cost \$250,000 Each and Falcon 9 Costs Less than \$30 Million. Available at <https://www.nextbigfuture.com/2019/12/spacex-starlink-satellites-cost-well-below-500000-each-and-falcon-9-launches-less-than-30-million.html> (2022/05/12) (2019)
23. Cardwell, Neal and Cheng, Yuchung and Gunn, C Stephen and Yeganeh, Soheil Hassas and Jacobson, Van: BBR: Congestion-based congestion control. *Queue* (2016)
24. Christopher S. Allen and Martina Giraudo and Claudio Moratto and Nobuyasu Yamaguchi: Chapter 4 - Spaceflight environment. In: Tommaso Sgobba and Barbara Kanki and Jean-François Clervoy and Gro Mjeldheim Sandal (ed.) *Space Safety and Human Performance*, pp. 87–138 (2018)
25. Daehnick, C., Klinghoffer, I., Maritz, B., , Wiseman, B.: Large LEO satellite constellations: Will it be different this time? Available at <https://www.mckinsey.com/industries/aerospace-and-defense/our-insights/large-leo-satellite-constellations-will-it-be-different-this-time> (2022/05/12) (2020)
26. del Portillo, I., Cameron, B.G., Crawley, E.F.: A technical comparison of three low earth orbit satellite constellation systems to provide global broadband. *Acta Astronautica* **159**, 123–135 (2019)
27. Dong, Mo and Li, Qingxi and Zarchy, Doron and Godfrey, Philip Brighten and Schapira, Michael: PCC: Re-architecting Congestion Control for Consistent High Performance. In: *Proc. of USENIX NSDI'15* (2015)
28. Dong, Mo and Meng, Tong and Zarchy, Doron and Arslan, Engin and Gilad, Yossi and Godfrey, Brighten and Schapira, Michael: PCC vivace: Online-learning congestion control. In: *Proc. of USENIX NSDI'18* (2018)
29. Duplyakin, D., Ricci, R., Maricq, A., Wong, G., Duerig, J., Eide, E., Stoller, L., Hibler, M., Johnson, D., Webb, K., Akella, A., Wang, K., Ricart, G., Landweber, L., Elliott, C., Zink, M., Cecchet, E., Kar, S., Mishra, P.: The design and operation of CloudLab. In:

- Proceedings of the USENIX Annual Technical Conference (ATC). pp. 1–14 (Jul 2019), <https://www.flux.utah.edu/paper/duplyakin-atc19>
30. Ekici, E., Akyildiz, I.F., Bender, M.D.: A distributed routing algorithm for datagram traffic in leo satellite networks. *IEEE/ACM Transactions on networking* **9**(2), 137–147 (2001)
 31. Emma Woollacott: Starlink Terminals Smuggled Into Iran - But How Effective Can They Be? Available at <https://www.forbes.com/sites/emmawoollacott/2022/10/25/starlink-terminals-smuggled-into-iran-but-how-effective-can-they-be/?sh=1c2952561027>(2022/10/28) (2022)
 32. Foust, J.: SpaceX adds laser crosslinks to polar Starlink satellites. Available at <https://spacenews.com/spacex-adds-laser-crosslinks-to-polar-starlink-satellites/> (2021/01/26) (2021)
 33. Giuliani, G., Klenze, T., Legner, M., Basin, D., Perrig, A., Singla, A.: Internet backbones in space. In: *ACM SIGCOMM Computer Communication Review*. vol. 50, pp. 25–37. ACM New York, NY, USA (2020)
 34. Giuliani, Giacomo and Ciussani, Tommaso and Perrig, Adrian and Singla, Ankit: ICARUS: Attacking low Earth orbit satellite networks. In: *Proc. of USENIX ATC '21*. pp. 317–331 (2021)
 35. Goyal, P., Agarwal, A., Netravali, R., Alizadeh, M., Balakrishnan, H.: ABC: A simple explicit congestion controller for wireless networks. In: *Proc. of USENIX NSDI '20* (2020)
 36. Goyal, P., Shah, P., Zhao, K., Nikolaidis, G., Alizadeh, M., Anderson, T.E.: Backpressure flow control. In: *Proc. of USENIX NSDI '22*. pp. 779–805 (2022)
 37. Handley, M.: Starlink revisions, Nov 2018. Available at https://www.youtube.com/watch?v=QEIUdMiColU&ab_channel=MarkHandley
 38. Handley, Mark: Delay is Not an Option: Low Latency Routing in Space. In: *Proc. of HotNets '18*. p. 85–91 (2018)
 39. Handley, Mark: Using Ground Relays for Low-Latency Wide-Area Routing in Megaconstellations. In: *Proc. of HotNets '19*. p. 125–132 (2019)
 40. Harris, M.: SpaceX Claims to Have Redesigned Its Starlink Satellites to Eliminate Casualty Risks. Available at <https://spectrum.ieee.org/spacex-claims-to-have-redesigned-its-starlink-satellites-to-eliminate-casualty-risks> (2019/03/21) (2019)
 41. Hauri, Y., Bhattacharjee, D., Grossmann, M., Singla, A.: "internet from space" without inter-satellite links. In: *Proc. of HotNets '20*. pp. 205–211 (2020)
 42. Henderson, T.R., Katz, R.H.: On distributed, geographic-based packet routing for leo satellite networks. In: *Proc. of GLOBECOM '00*. vol. 2, pp. 1119–1123. IEEE (2000)
 43. Hoots, F.R., Roehrich, R.L.: *Models For Propagation of Norad Element Sets* (1980)
 44. Hu, M., Xiao, M., Xu, W., Deng, T., Dong, Y., Peng, K.: Traffic engineering for software defined leo constellations. *IEEE Transactions on Network and Service Management* (2022)
 45. Hussein, M., Jakllari, G., Paillassa, B.: On routing for extending satellite service life in leo satellite networks. In: *Proc. of IEEE GLOBECOM '14*. pp. 2832–2837. IEEE (2014)
 46. Kassing, Simon and Bhattacharjee, Debopam and Águas, André Baptista and Saethre, Jens Eirik and Singla, Ankit: Exploring the “Internet from space” with Hypatia. In: *Proc. of ACM IMC '20*. pp. 214–229 (2020)
 47. Keith Cowing: Euroconsult report addresses challenges and potential of optical communications for nascent space applications market. Available at <https://spaceref.com/space-commerce/euroconsult-report-addresses-challenges-and-potential-of-optical-communications-for-nascent-space-applications-market/> (2023/01/31) (2023)

48. Kim Ancin: 5G + LEO: Verizon and Project Kuiper team up to develop connectivity solutions. Available at <https://www.verizon.com/about/news/5g-leo-verizon-project-kuiper-team> (2022/05/12) (2021)
49. Klenze, T., Giuliani, G., Pappas, C., Perrig, A., Basin, D.: Networking in heaven as on earth. In: Proc. of HotNets '18. pp. 22–28 (2018)
50. Kuiper USASAT-NGSO-8A ITU filing: USA2019-12905. Available at <https://www.itu.int/ITU-R/space/asreceived/Publication/DisplayPublication/8716> (2018)
51. Kuiper USASAT-NGSO-8B ITU filing: USA2019-13020. Available at <https://www.itu.int/ITU-R/space/asreceived/Publication/DisplayPublication/8774> (2018)
52. Kuiper USASAT-NGSO-8C ITU filing: USA2019-12909. Available at <https://www.itu.int/ITU-R/space/asreceived/Publication/DisplayPublication/8718> (2018)
53. LeoSat: Technical Overview. Available at <https://www.leosat.com/to/media/1114/leosat-technical-overview.pdf>
54. Li, Y., Li, H., Liu, L., Liu, W., Liu, J., Wu, J., Wu, Q., Liu, J., Lai, Z.: "internet in space" for terrestrial users via cyber-physical convergence. In: Proc. of HotNets '21. pp. 163–170 (2021)
55. Li, Y., Miao, R., Liu, H.H., Zhuang, Y., Feng, F., Tang, L., Cao, Z., Zhang, M., Kelly, F., Alizadeh, M., Yu, M.: Hpcc: High precision congestion control. In: Proc. of ACM SIGCOMM '19 (2019)
56. Lin, X., Rommer, S., Euler, S., Yavuz, E.A., Karlsson, R.S.: 5G from Space: An Overview of 3GPP Non-Terrestrial Networks. IEEE Communications Standards Magazine **5**(4), 147–153 (2021)
57. Ma, J., Qi, X., Liu, L.: An effective topology design based on leo/geo satellite networks. In: International Conference on Space Information Network. pp. 24–33. Springer (2017)
58. Micah Maidenberg, A.S.: Delta Air Lines Tested SpaceX's Starlink Internet for Planes Delta CEO Says. Available at <https://www.wsj.com/articles/delta-air-lines-tested-spacexs-starlink-internet-for-planes-delta-ceo-says-11650316287>(2022/10/28) (2022)
59. Michael Baylor: With Block 5, SpaceX to Increase Launch Cadence and Lower Prices. Available at <https://www.nasaspaceflight.com/2018/05/block-5-spacex-increase-launch-cadence-lower-prices/> (2022/05/12) (2018)
60. Mohorcic, M., Werner, M., Svigelj, A., Kandus, G.: Adaptive routing for packet-oriented intersatellite link networks: performance in various traffic scenarios. IEEE Transactions on Wireless Communications **1**(4), 808–818 (2002)
61. Netravali, Ravi and Sivaraman, Anirudh and Das, Somak and Goyal, Ameesh and Winstein, Keith and Mickens, James and Balakrishnan, Hari: Mahimahi: Accurate {Record-and-Replay} for {HTTP}. In: Proc. of USENIX ATC '15. pp. 417–429 (2015)
62. Niederstrasser, Carlos: Small Launch Vehicles-A 2018 State of the Industry Survey. In: Proc. of AIAA/USU Conference on Small Satellites (2018)
63. Papapetrou, E., Karapantazis, S., Pavlidou, F.N.: Distributed on-demand routing for leo satellite systems. Computer networks **51**(15), 4356–4376 (2007)
64. Perkins, C.E., Royer, E.M.: Ad-hoc on-demand distance vector routing. In: Proc. of IEEE Workshop on Mobile Computing Systems and Applications (WMCSA'99). pp. 90–100. IEEE (1999)
65. Plac, C.: LEO Speed: When Milliseconds are Worth \$Millions... An NSR Insight. Available at <https://news.satnews.com/2020/10/14/leo-speed-when-millisecond-s-are-worth-millions-an-nsr-insight/>(2022/10/31) (2020)
66. Schlicht, A., Marz, S., Stetter, M., Hugentobler, U., Schäfer, W.: Galileo pod using optical inter-satellite links: A simulation study. Advances in Space Research **66**(7), 1558–1570 (2020)

67. Siddiqi, A., Mellein, J., de Weck, O.: Optimal reconfigurations for increasing capacity of communication satellite constellations. In: 46th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference. p. 2065 (2005)
68. Sidibeh, K.: Adaption of the IEEE 802.11 protocol for inter-satellite links in LEO satellite networks. Ph.D. thesis, University of Surrey (United Kingdom) (2008)
69. Smutny, B., Kaempfer, H., Muehlhnikel, G., Sterr, U., Wandernoth, B., Heine, F., Hildebrand, U., Dallmann, D., Reinhardt, M., Freier, A., et al.: 5.6 Gbps optical intersatellite communication link. In: Free-space laser communication technologies XXI. vol. 7199, pp. 38–45. SPIE (2009)
70. SpaceX: FCC Selected Application for Space Exploration Holdings, LLC. Available at <https://fcc.report/IBFS/SAT-LOA-20161115-00118/1158350.pdf> (2016)
71. SpaceX: SPACEX NON-GEOSTATIONARY SATELLITE SYSTEM. Available at <https://fcc.report/IBFS/SAT-MOD-20190830-00087/1877671> (2019)
72. SpaceX FCC Filing: SpaceX V-BAND NON-GEOSTATIONARY SATELLITE SYSTEM. Available at https://licensing.fcc.gov/myibfs/download.do?attachment_key=1190019 (2017)
73. SpaceX Update: SPACEX NON-GEOSTATIONARY SATELLITE SYSTEM. Available at https://licensing.fcc.gov/myibfs/download.do?attachment_key=1569860 (2017)
74. SpaceX Update: SPACEX NON-GEOSTATIONARY SATELLITE SYSTEM. Available at <https://fcc.report/IBFS/SAT-MOD-20200417-00037> (2020)
75. Starlink: <https://www.starlink.com/legal/documents/DOC-1002-69942-69>. Available at <https://www.starlink.com/legal/documents/DOC-1002-69942-69> (2022/10/31) (2022)
76. Taleb, T., Mashimo, D., Jamalipour, A., Kato, N., Nemoto, Y.: Explicit load balancing technique for ngeo satellite ip networks with on-board processing capabilities. *IEEE/ACM transactions on Networking* **17**(1), 281–293 (2008)
77. Telesat: Telesat’s responses - Federal Communications Commission. Available at [http://licensing.fcc.gov/myibfs/download.do?attachment_key=1205775](https://licensing.fcc.gov/myibfs/download.do?attachment_key=1205775) (2018)
78. Telesat: APPLICATION FOR MODIFICATION OF MARKET ACCESS AUTHORIZATION. Available at <https://fcc.report/IBFS/SAT-MPL-20200526-00053/2378318.pdf> (2020)
79. Vivek Wadhwa and Alex Salkever: 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/> (2022/05/12) (2022)
80. Vladimirova, T., Sidibeh, K.: Inter-satellite links in leo constellations of small satellites (2007)
81. de Weck, O.L., de Neufville, R., Chaize, M.: Staged deployment of communications satellite constellations in low earth orbit. *Journal of Aerospace Computing, Information, and Communication* **1**(3), 119–136 (2004)
82. Yan, Francis Y and Ma, Jestin and Hill, Greg D and Raghavan, Deepti and Wahby, Riad S and Levis, Philip and Winstein, Keith: Pantheon: the training ground for Internet congestion-control research. In: Proc. of USENIX ATC ’18. pp. 731–743 (2018)
83. Young, C.: SpaceX’s Starlink internet will soon be available aboard cruise ships and airplanes. Available at <https://interestingengineering.com/innovation/spacex-s-starlink-internet-available-cruise-ships-airplanes> (2022/10/28) (2022)
84. Zech, H., Heine, F., Tröndle, D., Seel, S., Motzigemba, M., Meyer, R., Philipp-May, S.: Lct for edrs: Leo to geo optical communications at 1, 8 gbps between alphasat and sentinel 1a. In: Unmanned/Unattended Sensors and Sensor Networks XI; and Advanced Free-Space Optical Communication Techniques and Applications. vol. 9647, pp. 85–92. SPIE (2015)

85. Zech, H., Biller, P., Heine, F., Motzigemba, M.: Optical intersatellite links for navigation constellations. In: Sodnik, Z., Karafolas, N., Cugny, B. (eds.) International Conference on Space Optics (ICSO '18). vol. 11180, pp. 370 – 379 (2019)
86. Zhang, S., Li, X., Yeung, K.L.: Segment routing for traffic engineering and effective recovery in low-earth orbit satellite constellations. Digital Communications and Networks (2022)