

On the Feasibility of Laser Inter-satellite Links for Low-latency High Frequency Trading

Vaibhav Singh

A thesis

submitted in partial fulfillment of the
requirements for the degree of

Master of Science in Electrical Engineering

University of Washington
2023

Reading Committee:
Akshay Gadre, Chair
Payman Arabshahi

Program Authorized to Offer Degree:
Electrical & Computer Engineering

© Copyright 2023

Vaibhav Singh

University of Washington

Abstract

On the Feasibility of Laser Inter-satellite Links for Low-latency High Frequency Trading

Vaibhav Singh

Chair of the Supervisory Committee:

Professor Akshay Gadre
Electrical & Computer Engineering

Internet satellite constellations are proliferating and enabling new applications in long-range and large-scale connectivity. One such promising application is providing low-latency communication for high frequency trading by leveraging their global coverage. While existing approaches leverage higher orbit satellites as low-latency relays for private clients, there is a new opportunity that can provide low-latency connectivity to every customer via the satellite mesh – laser inter-satellite links (LISLs).

While past research has explored the capability of LISLs for enabling physical deployment and impact on the constellation topology design, this work presents a holistic empirical study on the feasibility of leveraging LISLs for low-latency high frequency trading. We evaluate the feasibility, utility and practical limits of the latency benefits of LISL-based constellations for ten large and geographically diverse financial markets across over 150 million satellite-seconds of simulation. Our empirical evaluation across two of the largest proposed satellite constellations, Starlink and Project Kuiper, demonstrates that LISLs can provide full coverage with roughly 1000 satellites and improve latency for 78.76% of link-pairs. However, these benefits decline drastically with processing delays of 4 ms to roughly 19.05% of link-pairs.

Acknowledgements

First and foremost, I want to express my sincere gratitude to my advisor, Professor Akshay Gadre for his invaluable advice, feedback and undeterred patience that played an immeasurable role in shaping the trajectory of my research and its ultimate success. Your profound insights, shared during our numerous discussions and meetings encouraged me to dig in and explore previously unanswered questions. Thank you for believing in me and your endless support! I would also like to extend many thanks to Professor Payman Arabshahi for accepting my committee invitation, providing his encouragement and mentorship, as well as enriching my graduate school journey by helping me grow as an instructor.

Special mention to Shengzhi for giving me a crash course on Keynote and my roommate Pushkar for his constant support! Thank you to everyone at the NSL lab and NEWTLab for giving me honest feedback and brainstorming approaches to dissect problem statements! I also thank my graduate advisor Jennifer Huberman for helping me navigate this graduate school journey.

Contents

1	Introduction	6
2	Related Work	10
2.1	Satellite Constellations and Topologies	10
2.2	Laser Inter-satellite Links (LISLs)	11
2.3	Low-latency Satellite Relay	11
3	Motivation & Background	12
3.1	High Frequency Trading	12
3.2	Satellite Constellations	13
4	Objectives & Approach	17
4.1	Objectives	17
4.2	Approach	18
5	Evaluation	21
5.1	Feasibility	22
5.2	Utility	24
5.3	Practical Limits	26
6	Discussion	30
6.1	Take-aways	32
7	Conclusion	33
8	Future Work	34

List of Figures

1.1	Traditional Satellite Constellations function as a relay, relying on existing infrastructure to provide internet connectivity. LISL-based Satellite Constellations have the potential to behave as a full-fledged backhaul.	7
3.1	Walker-Delta Patterns and LISLs: A satellite constellation can be modeled as a Walker-Delta pattern. Permanent LISLs can only be formed between nodes with small relative velocities.	14
4.1	LISL Evaluation Framework: The evaluation framework performs trace-driven emulation of the satellite constellations to empirically measure the latency benefit provided by LISL-based backhauls in the presence of practical constraints.	19
5.1	LISL Feasibility (a) Coverage vs. # of satellites: LISLs can achieve complete coverage across the financial hubs with both Kuiper and Starlink meshes with roughly 1000 satellites.	22
5.2	LISL Feasibility (b) Impact of Phasing Parameter on coverage: The right choice of the phasing parameter F can provide significant coverage benefits for LISL based satellite constellations (Starlink).	23
5.3	LISL Feasibility (c) Maximum Delay between coverage across links: LISLs in Kuiper constellations will face lower coverage unavailability periods compared to Starlink constellation due to the smaller inclination angle but will provide less coverage in polar regions.	24

5.4 **LISL Utility (a) Median LISL Latency benefit across Pairs:** Across time, 27.11% of LISL based constellations can achieve half the terrestrial latency, about half the links see 37% lower latency and 78.76% links see a benefit. 24

5.5 **LISL Utility (b) Median Latency across Pairs:** LISL-based constellations can provide significant latency benefits based on propagation delay. However, terrestrially colocated locations face comparable or lower latency via ground infrastructure. Poorly interconnected locations, especially in southern hemisphere, achieve the maximum benefit via LISLs. 25

5.6 **LISL Utility (c): Benefit vs. Ground Latency:** While ground links face geographical constraints, constellations can circumvent these obstacles to provide lower latency connectivity. 26

5.7 **Practical Limits (a) Processing delay Impact:** The efficacy of LISL-based constellations to provide latency benefits drops significantly with increasing values of processing delays. 27

5.8 **Practical Limits (b) Impact of geographical locations:** Similar to terrestrial links, constellations face relatively higher median latency for connectivity across hemispheres which will be significantly increased with processing delay. 28

5.9 **Practical Limits (c) Impact of terminals on latency benefit:** In the presence of processing delay, NOP links show a definitive latency advantage over just IOP + AOP LISLs. 28

6.1 Kuiper is expected to deploy three satellite meshes with different altitudes and inclination angles in its constellation which will not be able to form permanent LISLs across them. 31

Chapter 1

Introduction

Internet satellite constellations are becoming an increasingly viable alternative to ground-based Internet connections, especially in rural and underserved areas [1]. These constellations consist of thousands of satellites that have the potential to form an interconnected mesh around the world and beam down internet connectivity to terrestrial base stations. One promising application for leveraging these low-latency meshes enabled by these inter-satellite links is in high-frequency trading [2]. Indeed, there have been significant efforts by the largest companies in the satellite constellation space to develop a special category of service for high-throughput low-latency communication [3]. While today's satellites provide low-latency communication by relaying through dedicated ground stations deployed in select locations and providing priority to these network flows, the true frontier for HFT lies in leveraging laser inter-satellite links (LISLs) to achieve better latency than terrestrial networks.

The reason LISLs may provide better latency than terrestrial networks is three-fold: (1) Signal propagation is $\sim 47\%$ faster over air/vacuum than in optical fiber cables; (2) Trans-continental fiber-optic connections do not form a straight path between the largest financial markets; (3) Terrestrial networks can face queuing, processing and firewall bottlenecks that amplify latency across continents. Given the recency in development of LISL prototypes, much of the related work has primarily focused on understanding impact of constellation topology design [4, 5], characterizing practical limits [6] and establishing individual inter-satellite links [7, 8]. This work aims to fill the

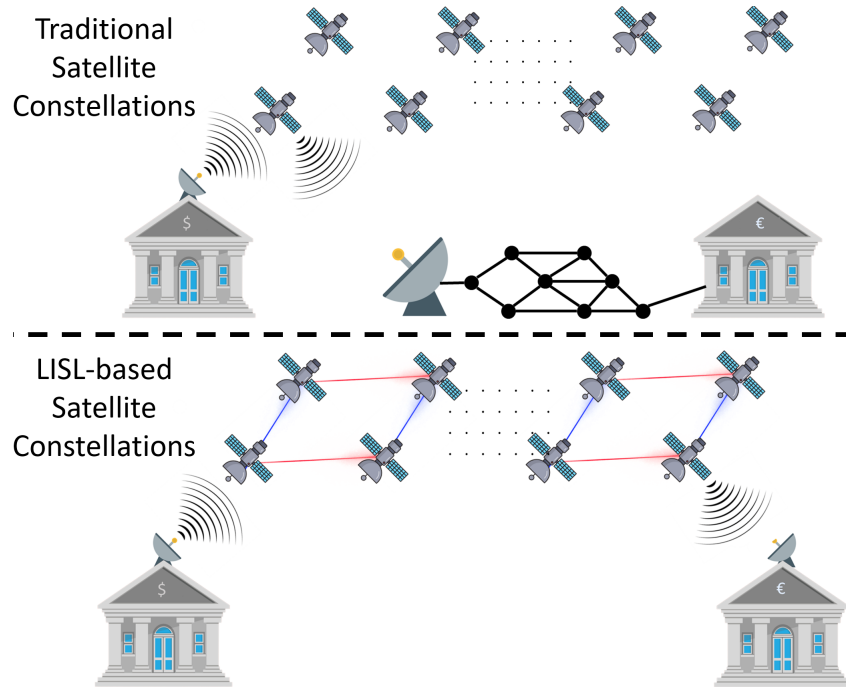


Figure 1.1: Traditional Satellite Constellations function as a relay, relying on existing infrastructure to provide internet connectivity. LISL-based Satellite Constellations have the potential to behave as a full-fledged backhaul.

gap where there is no dedicated empirical study that characterizes the feasibility and amount of latency benefit provided by such LISL-based constellations for high frequency trading.

This work presents a holistic evaluation of the latency benefit provided by leveraging LISL-based constellations for high frequency trading. To achieve this, we focused on the ten large and geographically distributed financial hubs¹ and analyze the feasibility of leveraging LISLs for communicating high frequency trading transactions. Further, instead of focusing on aggregate measurements, it is useful to identify the critical limits of the various factors, such as processing delay and geographical locations, that impact the amount of latency benefit provided by LISL-based backhauls. We evaluate the above metrics by performing trace-driven evaluation on the satellite meshes of the two largest existing and proposed satellite constellations, Starlink and Project Kuiper. This study measures the three metrics as described below:

¹New York (NYSE/NASDAQ), London (LSE), Hong Kong (HKSE), Mumbai (BSE/NSE), Frankfurt (DBAG), Paris (Euronext), Sydney (ASE), Johannesburg (JSE), São Paulo (B3 S.A.), Tokyo (JEG)

Feasibility (*Can LISL constellations provide connectivity across geographically diverse end-nodes?*): There are several factors that govern the coverage of a satellite constellation across geographically diverse locations. We focus on how the coverage changes as more satellites are launched periodically to form a new constellation. Next, we study how these new LISL-based satellite constellations can be optimized for the latency benefit coverage by identifying optimal values of constellation parameters such as the phasing parameter. The evaluation and results are described in Sec. 5.1.

Utility (*Can LISL constellations provide lower latency over traditional networks?*): Given the coverage availability of potential satellite constellations, it is important to analyze the benefit of using this LISL mesh for achieving lower latency communication over alternatives. Comparison of current terrestrial latency with the satellite mesh latency is done. Although LISLs benefit from $\sim 47\%$ faster links compared to ground networks, factors such as added altitude and larger mesh radius necessitate a deeper analysis of the end-to-end latency benefit. The evaluation and results are described in Sec. 5.2.

Practical Limits (*What are the impacts of practical considerations on the benefits of LISLs?*): There are several practical limitations that may limit the ability of LISLs to provide lower latency communication, such as processing and queuing delays as well as the kind of terminals deployed on the satellites [6]. This work examines these important considerations inherent in relaying information across LISLs. The analysis focuses on the stringent limits that these delays add to the latency and the impact of the choice of terminals while designing these satellite constellations. The insights are in Sec. 5.3.

The evaluation framework was implemented in MATLAB and measured the above metrics by emulating two well-known proposed constellations, SpaceX’s Starlink and Amazon’s Project Kuiper. Their Walker-Delta constellations were modeled, operating in perfect meshes across 157.3 million satellite-seconds of simulation for the 45 link-pairs. The most interesting findings are:

- Both Starlink and Kuiper constellations with LISL-based backhauls can provide complete coverage across the link-pairs with 1000 satellites.
- 78.76% of all link-pairs across time can achieve a latency improvement using LISLs with a median latency improvement of 37%. 27.11% of link-pairs can achieve half their terrestrial latency using LISL-based constellations.
- The number of link-pairs benefiting from LISLs falls drastically to 40.5% and 19.05% with small processing delays of 2 ms and 4 ms per satellite, respectively.

Based on the above results, the following recommendations shall be of use to researchers and industry:

- *New satellite constellation deployments should put emphasis on choosing the right values of the phasing parameter to provide better coverage.*
- *LISL latency benefits are more pronounced between ground station pairs where the path through terrestrial infrastructure is longer than the orthodromic distance.*
- *LISL network developers should prioritize reducing processing and queuing delays on satellites to below a millisecond to achieve significant latency benefits.*

Chapter 2

Related Work

This section highlights some of the related work that has explored similar problems in relevant areas:

2.1 Satellite Constellations and Topologies

While older satellite constellations operate at much higher altitudes (GEO/MEO) [9], the recent reduction in cost of satellite launches has opened new opportunities for exploring satellite constellations at much lower orbits. Low Earth Orbit (LEO) satellites have much lower round-trip times compared to geostationary satellites, but require a large number of satellites to provide continuous global coverage [10].

In the past decade, several firms have entered the race to deploy internet satellite constellations including Telesat's LightSpeed [11, 12], Viasat's ViaSat-3 [13], SpaceX's Starlink [14] and Amazon's Project Kuiper [15]¹. While some such as Starlink provide services direct-to-customer, others including Telesat provide services to large business clients such as commercial airliners. These satellite constellations consist of hundreds to thousands of LEO satellites with the aim of providing continuous coverage across the globe. Such constellations can enable remote-sensing systems and Internet-of-Things (IoT) without broadband access [16]. They can also enable lower latency communication via ground-satellite cooperation [17]. Other related works have also explored optimizing

¹Kuiper is used as a proxy for Project Kuiper for the rest of this text.

satellite constellation parameters to improve coverage [18, 19, 20] and avoid collisions [21]. In contrast, this work does not focus on optimizing satellite constellations for enabling LISL-based HFT, but instead aims to study the impact of design decisions made by the above constellation designers on the feasibility of low-latency communication via LISL backhauls.

2.2 Laser Inter-satellite Links (LISLs)

Researchers have already demonstrated laser links between satellites. For example, the European Data Relay System (EDRS) [7, 22] provides a 1.8 Gbps link between the system’s geostationary satellites and LEO client satellites as a relay. Early stage testing of laser terminals has also been performed with promising results [23]. Recent studies have also analyzed various factors that govern the formation and operation of LISLs such as link establishment delay [8], network metrics for constellations [24] and dynamic routing across temporary links [25, 4, 26].

More recently, researchers have theorized leveraging permanent LISLs for creating a mesh between constellation satellites [5, 6] as well as their behavior across constellations [27]. This study takes a more holistic end-to-end approach to evaluate the latency benefits and practical limits of these links across today’s proposed constellations.

2.3 Low-latency Satellite Relay

Leveraging satellites to perform low-latency communication has been explored in various ways from providing hybrid low-latency communication [28] as well as adapting constellation connectivity to achieve low latency [29]. The use of skywave in shortwave spectrum for inter-continental distances is gaining traction, but faces stringent opposition from the enthusiast ham radio community [30]. Iridium [31] and Iridium NEXT [32] have also provided satellite telephony services for users. Further, the impact of hybrid topology on network latency and jitter has been analyzed [33]. In contrast, we consider a satellite constellation network where the HFT data is completely routed using LISLs between the satellites and measure its ability to provide latency benefits compared to terrestrial networks.

Chapter 3

Motivation & Background

In this section, we provide the primary motivation of leveraging LISLs for low-latency communication as well as background on satellite deployments today.

3.1 High Frequency Trading

High frequency trading (HFT) is an approach where investors, hedge funds, and more recently, data-driven algorithms use complex strategies to profit from small variations in share and options prices. Due to the need to make quick and judicious decisions, HFT companies can see enormous benefit from even small decreases in latency [34]. By one estimate, a 1-millisecond advantage can translate to \$100 million a year of profit to a major brokerage firm [35]. Furthermore, HFT is thought to account for more than half of all US equity trades [36]. This requirement of low latency compels a large number of HFT firms to reduce latency by co-locating near financial markets [37], and devising private high-priority networking contracts, sometimes making use of microwave networks [38]. Companies have even gone so far as to designing dedicated FPGAs to facilitate faster execution of trading decisions [39]. However, despite these high-cost approaches, there still remains a fundamental bottleneck in achieving lower latency – speed of light in optical fiber.

In fact, just the propagation delay of a signal traveling across several hundred kilometers of optical fiber can accumulate to several milliseconds. Furthermore, queuing and processing delays present in terrestrial networks due to other network flows can make it difficult to reduce latency

beyond a certain point. Recently, satellite relays have been thought of as a solution to this final bottleneck of latency [2]. However, while such relays may work for niche deployments, the true solution en-mass is expected to be enabled by a mesh of Low Earth Orbit (LEO) satellite constellations inter-connected with laser inter-satellite links (LISLs).

3.2 Satellite Constellations

Several major projects in the satellite industry such as Starlink and Kuiper have proposed plans to provide connectivity to terrestrial users via LEO satellite constellations. Today's constellation deployments operate by acting as communication relays that forward the packets between the user and a designated terrestrial ground station connected to the internet. However, the long-term goal for these constellation deployments is to eventually operate as a cohesive backhaul for providing connectivity across Earth.

In its most basic form, a satellite constellation mesh can be characterized by a Walker-Delta Constellation pattern [40] which describes a constellation of uniformly spaced satellites in both latitude and longitude. Walker-Delta pattern can be used to describe a constellation with the following parameters (see Fig. 3.1):

Inclination Angle (ϕ): This is the angle made by the satellite paths with the equatorial plane of the earth as they traverse across their orbits. While polar satellites have an inclination angle closer to 90° , most constellation deployments prefer inclination angles between 30° to 60° for optimizing cost and coverage.

Number of Geometry Planes (M): This is the distribution of the number of satellites in the constellations across the longitudinal axis. An increase in the number of geometry planes improves the instantaneous longitudinal coverage of the constellation while a decrease translates to better instantaneous latitudinal coverage. Traditionally, constellations are deployed to equalize the number of geometry planes and the number of satellites in those planes to achieve optimal distribution of coverage.




Figure 3.1: Walker-Delta Patterns and LISLs: A satellite constellation can be modeled as a Walker-Delta pattern. Permanent LISLs can only be formed between nodes with small relative velocities.

Phasing Parameter (F): This parameter defines the relative offset between two closest satellites in adjacent orbital planes. Typically, F can be varied between 0 and M and this distribution translates to the adjacent satellite being offset to roughly the next satellite in its path. One of the more surprising consequences of varying F is the fact that the spacing between satellites of two adjacent planes changes across time and affects the coverage of the constellation. Furthermore, the choice of F plays an important factor in avoiding collisions between satellites of the same constellation operating at the same altitude.

Total Number of Satellites (T): This parameter defines the total number of satellites in the constellations that are distributed across the M geometry planes. As can be expected, increasing the number of satellites also increases coverage of a constellation across the Earth.

A short-form representation of satellite constellations can be denoted as a Walker-Delta 4-tuple $WD(\phi/T/M/F)$. However, given the large distances between satellites and the expectation of global coverage, both Starlink and Kuiper are planning to operate multiple constellation meshes in tandem with different altitudes, number of satellites and inclination angles. For example, based

on the FCC filing on Kuiper [15], it is expected to operate at least 3 constellations described by WD(42°/1512/42/F) at a 610 km altitude¹, WD(52°/1156/34/F) at a 630 km altitude and WD(33°/784/28/F) at a 590 km altitude. While these constellations with different operating altitudes will improve coverage and latency, it is quite likely that permanent LISLs will primarily be formed across satellites operating at the same altitude. (More analysis and discussion on multiple meshes in Sec. 6)

Laser Inter-Satellite Links: LISLs are communication links between two satellites established using a highly directed laser beam [41]. Using free space optics over traditional Radio Frequency (RF) links has several advantages in LEO. The key among them being lower transmit power requirements for the same received power at a given distance [42, 43]. This is crucial in an energy constrained environment such as satellites in space. LISLs also overcome the need for hopping network traffic between ground-stations and satellites when connecting far-off clients on Earth, drastically improving network latency. Instead, LISLs can enable a high-fidelity and high-throughput backhaul to provide mesh-wide connectivity within a constellation. However, the usage of lasers to connect two satellites is a double-edged sword – whereas a traditional RF link would provide connectivity to any satellite in line-of-sight (provided the received signal strength is greater than receiver sensitivity), optical laser connections can only create one-to-one connections. Current generation of Acquisition, Tracking and Pointing (ATP) systems are not capable of maintaining temporary links (i.e., links where the two satellites have high relative velocities) [44].

Much work has been done to design laser terminals and receivers, to develop the requisite technology for LISLs. This includes recent testing evaluation in space conditions [23]. However, the nascent nature of the technology makes much of this discussion rely on the scientific hypothesis of other scholarly work in this area. Given the constraint of forming links between satellites with minimal relative velocities, there have been three types of proposed links [5]:

Intra-orbital plane (IOP) links are established between satellites in the same orbital plane. While the satellites do face minimal relative velocities across an orbit, high fidelity links are feasible.

¹There are various reports that mention 1296 satellites with 36 geometry planes, however, this study follows the official FCC documentation.

Adjacent-orbital plane (AOP) links are formed between satellites in the adjacent geometry planes of the satellite constellations. These links face lower relative velocities closer to the equator and relatively higher relative velocities closer to the poles.

Nearby-orbital plane (NOP) links are developed between satellites in the non-adjacent geometry planes of the constellation. While these links face higher relative velocities, they also require less number of hops to traverse large distances.

While there have been other proposed link types such as Crossing-orbital plane (COP) links, these links face significantly higher relative velocities and remain infeasible in near future [44]. Further, most proposed satellites are unable to form more than two pairs of LISLs [45] which means there is an important design decision between the above links to optimize for latency, coverage and number of intermediate nodes in the LISL constellation mesh.

Chapter 4

Objectives & Approach

4.1 Objectives

Our goal in this study is to present a micro-analysis on the feasibility of leveraging LISLs for low-latency high frequency trading. We focus less on the constellation properties and more on the end-to-end link experience for users performing HFT transactions. Towards this goal, we pose three important research questions:

(1) Can a consistent LISL-based connection be formed for HFT transactions between link-pairs? There are various parameters as defined above that affect the feasibility of link formation such as satellite receivers, acquisition delays, and availability of bandwidth in the network mesh. However, many of these factors will be highly deployment-specific and governed by the individual stakeholders. Thus, this study focuses on three parameters that are independent of the deployment and are limited by physical constraints: (a) impact of the number of satellites (T) in the constellation; (b) impact of the phasing parameter (F); and (c) the length of time that coverage is unavailable to a given link pair. Note that much of the analysis will focus on ten large and geographically distributed financial centers, as mentioned previously.

(2) Does the LISL-based link lower latency compared to terrestrial links for HFT transactions? The objective is to ascertain the latency of propagation delay throughout the LISL-based mesh in the constellation and compare it with the well-documented terrestrial latency data

from Global Ping Statistics [46]. Further, given the organic deviation caused by movement of the constellation relative to ground station locations and its impact on latency, we identify the benefit achieved by comparing the median, 75%ile and 99%ile LISL latency with stable terrestrial median latency.

(3) How much do the practical real-world factors affect the benefit of LISL-based links?

We hope to identify the impact of factors that have a significantly large impact on terrestrial links such as geographical diversity, processing and queuing delays, and number of simultaneous links. Analogous to these factors, LISLs also face the problem of geographical diversity across longitudes and latitudes based on the phasing angle and distribution of the total number of satellites across geometry planes. Additionally, processing delay will be larger for these new technologies due to the sheer amount of data that will be communicated using these laser-based links [47]. Finally, most advanced satellites deployed today can only deploy 4 laser terminals with fixed directions [45]. The objective is to measure the impact of these parameters and characterize the goals for researchers optimizing LISL-based networks for HFT.

4.2 Approach

A first-principle approach was taken to perform an empirical analysis of the feasibility, utility, and practical limits of the LISL constellations in the context of HFT. Unlike a theoretical approach where it would be difficult to characterize several important temporal factors such as the oblate spheroid shape of earth, intrinsic wobble of earth, and the phasing parameter behavior, an empirical approach can leverage the models based on extensive data collected by prior researchers such as SGP4 [48] and SDP4 [49]. Further, the temporal features are quite important given that the relative motion of the satellite mesh to the stationary financial hub nodes can significantly affect both latency and connectivity.

Next, given the empirical approach, we evaluated the system across different time frames and step sizes to measure two metrics - Connectivity (*Can a LISL-based path be formed?*) and Latency (*What is the minimum latency?*). We then measured the aggregate statistics such as mean, median, the histogram distribution and evaluated the optimal parameters for performing this evaluation.

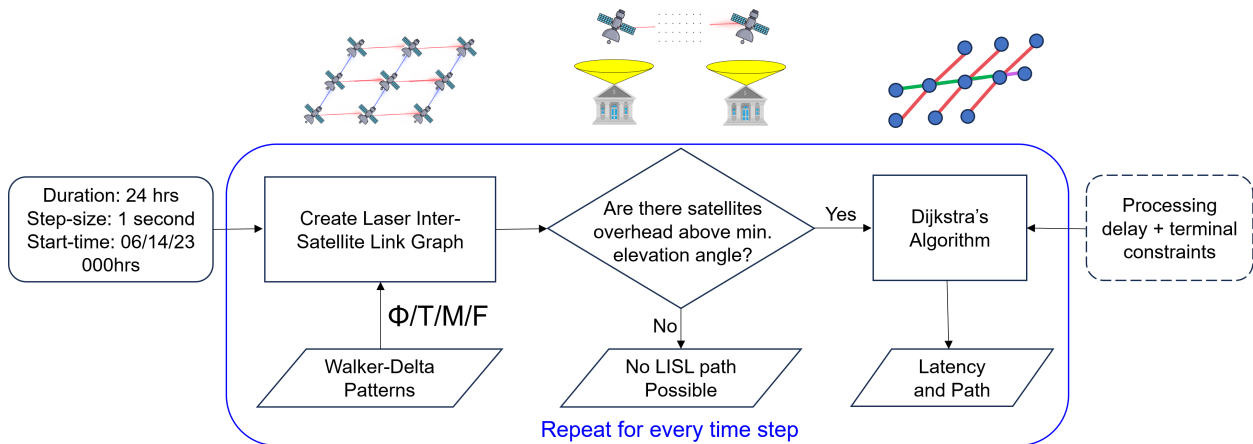


Figure 4.1: LISL Evaluation Framework: The evaluation framework performs trace-driven emulation of the satellite constellations to empirically measure the latency benefit provided by LISL-based backhuls in the presence of practical constraints.

The microbenchmark results showed that evaluation across 24 hours for every second can measure all of the required parameters capturing all facets of the two metrics effectively. We choose the 24-hours of the trace-driven simulation to happen on June 14, 2023 when the TLE files retrieved and Walker-Delta patterns were developed.

We choose the following two constellations based on Starlink TLE files [50] and Project Kuiper FCC filings [15] and abstract them as the following Walker-Delta patterns: Starlink Phase I Version 3 Constellation as $WD(53^\circ/1584/72/F)$ at an altitude of 550 km and Kuiper Constellation as $WD(42^\circ/1512/42/F)$ at an altitude of 610 km. Note that the phasing parameters (F) for both constellations are not publicly known and the same value is used across the evaluation to compare results. The impact and variation of this phasing parameter is much lower when the number of satellites is larger than thousand not affecting the sanctity of the results. We also choose ten ground stations locations (described in Sec. 5) based on their financial market-cap and geographical location diversity. Thereafter, connections were formed between these links to form a graph as follows:

- First connect the ground stations to the nearest satellite which is above the minimum elevation angle (typically 30° [51]) as an RF link (as it occurs today)
- Next, connect IOP links in every geometry plane of the constellation. This link provides

latitudinal coverage.

- Create two copies of this graph and form AOP and NOP links across geometry planes. These links provide longitudinal coverage.

Each of the above links are given the weight of the propagation delay in the graph with LOS constraints [5]. Further, the two graphs are created separately since real deployments can only support 4 laser terminals [45] per satellite today. Further, additional processing delay inherent in relaying information across laser terminals was modeled. While most terrestrial routers face a processing delay of several hundred microseconds, these lasers will likely forward several magnitude larger data resulting in larger processing delays.

Finally, we used Dijkstra’s algorithm [52] on the graph to find the lowest latency path through the constellation and measure both the number of hops and the minimum latency. Note that without a processing delay, the minimum latency path will always choose AOP links over NOP links, while in the presence of processing delay NOP links provide a route with less number of hops.

The framework was run for 24 hours starting at June 14th, 2023, 00:00 hrs GMT with a time step-size of 1 second to recalculate the orbital positions for the entire constellation. Across all the experiments, this roughly corresponds to more than 150 million satellite seconds of model-driven simulation.

Chapter 5

Evaluation

Given that LISLs are still quite nascent with little academic measurement and evaluation studies, it is unfeasible to perform the evaluation on a real scenario¹. Thus, we implemented the trace-based simulation framework in MATLAB and recreated the existing/proposed satellite constellation topology of Starlink and Kuiper. While both Starlink and Kuiper have multiple meshes to improve coverage and connectivity, each of these individual meshes follow quite closely with a Walker-Delta pattern as described in Sec. 3. Thus, our baseline comparison for these proprietary constellations across various parameters relies on the Walker-Delta patterns $WD(42^\circ/1512/42/F)$ for Kuiper and $WD(53^\circ/1584/72/F)$. Note that across our evaluation and comparison with the baseline, we preserve the number of satellite planes as it is pivotal in determining the success of the constellation topology. For the actual topologies, we extracted the publicly available TLE files for Starlink [50] and public FCC filings from Project Kuiper [15]. The framework simulates the number of satellites (varied across experiments and constellations described below) for every second across 24 hours, measures the shortest path latency and stores the end-to-end LISL path taken for the HFT transaction. As a baseline, we compare these latency estimates with ground latency measurements obtained from Global Ping Statistics [46] across a month between the locations.

We chose ten large financial stock exchange locations as a reasonable proxy for understanding the feasibility across the globe for performing high frequency trading via LISLs. We also emphasized

¹In fact, neither Starlink nor Project Kuiper has yet successfully publicly demonstrated operational LISLs between satellites

the importance of geographical and continental diversity in our evaluation to include several stock exchanges in the southern hemisphere, instead of purely focusing on the market-caps. I believe this choice provides a more equitable understanding of the limits of such LISLs for high frequency trading. The ten stock exchanges chosen based on the above criteria are: New York Stock Exchange (NY), London Stock Exchange (LD), Hong Kong Stock Exchange (HK), Bombay Stock Exchange (MU), Deutsche Börse AG (FR), Euronext (PA), Australian Securities Exchange (SY), Johannesburg Stock Exchange (JO), B3 Brasil Bolsa Balcão (SP), and Japan Exchange Group (TY).

5.1 Feasibility

In this section, we evaluate the feasibility of providing coverage to the links between these stock exchanges by leveraging LISL-based interconnects.

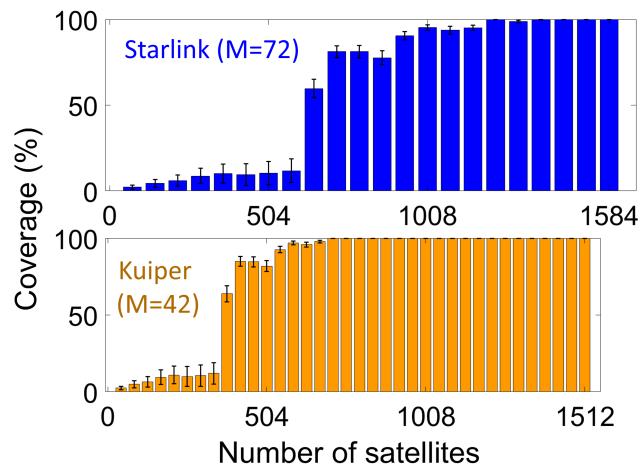


Figure 5.1: LISL Feasibility (a) Coverage vs. # of satellites: LISLs can achieve complete coverage across the financial hubs with both Kuiper and Starlink meshes with roughly 1000 satellites.

Coverage vs. # of satellites: The number of satellites in the Starlink and Kuiper constellations were increased while maintaining the original number of geometry planes and inclination angle. We evaluate whether the constellation can enable LISL-based link formation every second for 24 hours. Note that this only verifies the possibility of link formation and not necessarily the benefit over terrestrial links. The results described in Fig.5.1 show that beyond roughly 1000 satellites, the coverage for the formation of links is not bottlenecked by the number of satellites. Notice how

the Starlink constellation with a larger number of geometry planes requires a much larger number of satellites to provide latitudinal coverage but can achieve better coverage in regions at higher latitudes.

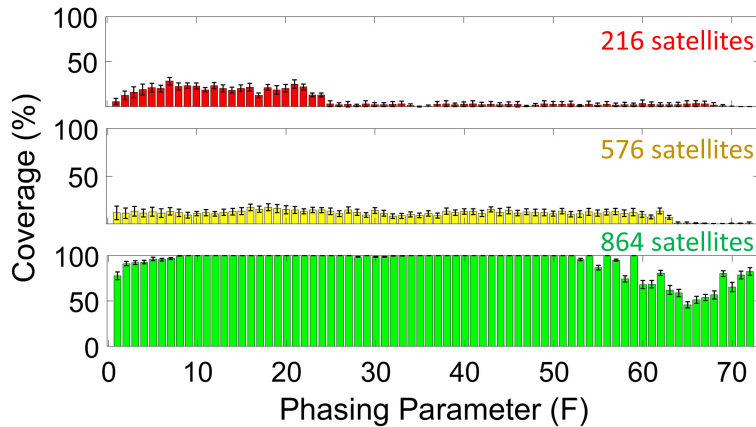


Figure 5.2: LISL Feasibility (b) Impact of Phasing Parameter on coverage: The right choice of the phasing parameter F can provide significant coverage benefits for LISL based satellite constellations (Starlink).

Impact of Phasing Parameter on coverage: In scenarios where there is not sufficient coverage, we evaluate the impact of the choice of the phasing parameter, F , on the coverage provided by the satellites of the Starlink constellation. The right choice of the phasing parameter F for providing coverage will maximize the distance between the closest satellites to optimize the coverage spread. The results in Fig.5.2 shows that the phasing parameter has a significant impact (as large as two-fold difference) on the coverage of the constellations with less than 1000 satellites. Further, while the choice of F changes as the number of satellites deployed increases, given the elevation angle, there exists a sweet spot (in our scenarios, 5 to 15) where across the number of satellites, the coverage is maximized. However, while determining the optimal phasing parameter, existing satellites at the same altitude must be respected in order to avoid collisions [53].

Maximum Delay between coverage across links: Finally, we evaluate the maximum delay between coverage across time as the number of satellites increases for both constellations. As shown in Fig. 5.3, with the lower inclination angle and less number of geometry planes, Kuiper achieves better coverage with fewer satellites. However, this comes with the important caveat of providing less coverage across the polar regions. Despite this, the median coverage duration loss is ~ 7 minutes

each time the link loses communication access through LISL-based backhauls. Thus, to minimize the maximum delay and coverage of these links, it is critical to increase the number of satellites deployed or derive a hybrid model with terrestrial links to provide coverage [17].

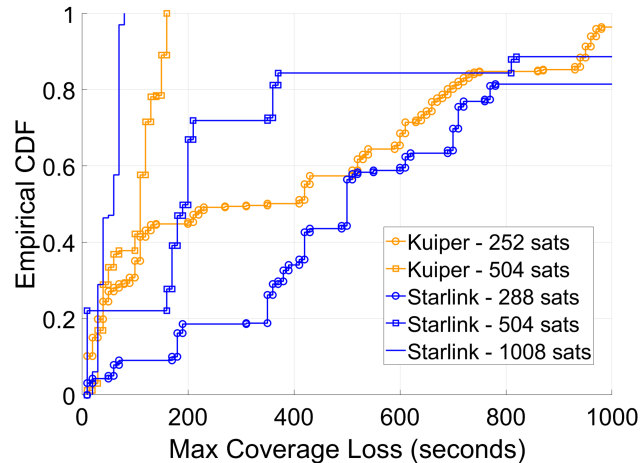


Figure 5.3: LISL Feasibility (c) Maximum Delay between coverage across links: LISLs in Kuiper constellations will face lower coverage unavailability periods compared to Starlink constellation due to the smaller inclination angle but will provide less coverage in polar regions.

5.2 Utility

In this section, we evaluate the utility of providing coverage to the links between these stock exchanges by leveraging LISL-based interconnects.




Figure 5.4: LISL Utility (a) Median LISL Latency benefit across Pairs: Across time, 27.11% of LISL based constellations can achieve half the terrestrial latency, about half the links see 37% lower latency and 78.76% links see a benefit.

LISL latency benefit across link-pairs: We evaluate and compare Starlink and Kuiper constellations as well as terrestrial ground links in terms of end-to-end latency achieved purely via propagation delays. Fig. 5.4 shows the median and quartile latency of all of these above links across the 45 links between the stock exchanges. The results demonstrate that LISL-based backhauls can provide significant benefits in terms of latency over ground links for geographically challenged locations where very few ground links exist (such as locations in the southern hemisphere, Sydney (SY), Johannesburg (JO) and São Paulo (SP)). In locations where terrestrial links are quite direct or close in orthodromic distance [54] such as between Paris (PA), Frankfurt (FR) and London (LD), terrestrial networks achieve lower latency compared to LISL-based satellite constellations. Finally, the remaining links face a much lower median latency via LISL-based constellations but comparable 90%ile latency to terrestrial links.

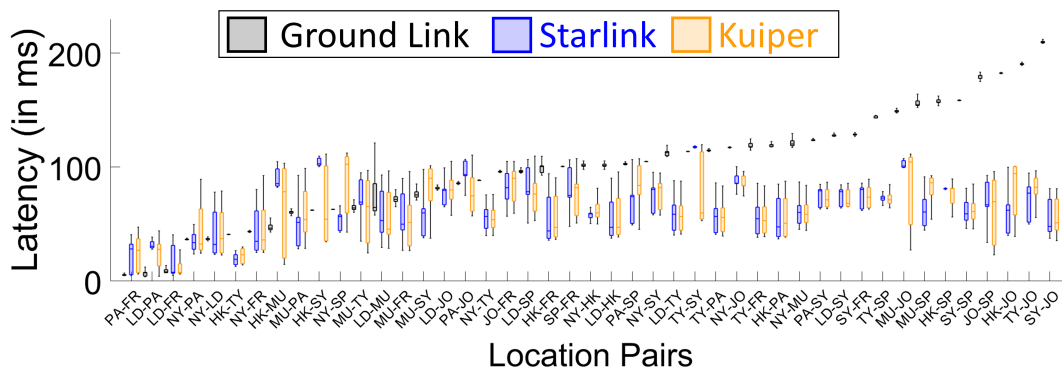


Figure 5.5: LISL Utility (b) Median Latency across Pairs: LISL-based constellations can provide significant latency benefits based on propagation delay. However, terrestrially colocated locations face comparable or lower latency via ground infrastructure. Poorly interconnected locations, especially in southern hemisphere, achieve the maximum benefit via LISLs.

Magnitude of LISL latency benefit: We next study the amount of improvement in latency that LISL-based satellite constellations can provide compared to terrestrial ground links. As shown in Fig. 5.5, about 27.11% of LISL pairs can achieve half the median ground latency. 50% of LISL pairs achieve roughly 37% lower latency via LISL satellite constellations compared to ground links. Finally, 78.76% of satellite links achieve better latency than the ground link would provide across time. Overall, given the constraint of only considering the propagation delay, LISL-based networks

can provide high value for HFT applications across large distances.

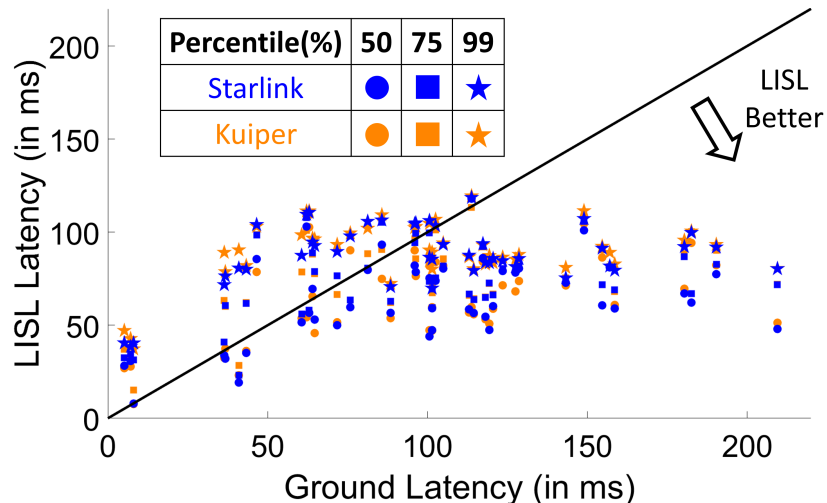


Figure 5.6: LISL Utility (c): Benefit vs. Ground Latency: While ground links face geographical constraints, constellations can circumvent these obstacles to provide lower latency connectivity.

Comparison of LISL latency with ground latency: We next study the impact of the ground distance (a good proxy for terrestrial latency) on the latency benefits of leveraging LISLs. The results in Fig. 5.6 show that LISLs would provide better coverage across longer distances due to the increased distance over which the latency of a link is computed. In fact, the result clearly indicates that there is an upper bound for LISL-constellation latency of ~ 120 ms while the median ground latency can be as high as 220 ms between Sydney and Johannesburg. However, this trend will curve upwards upon consideration of processing delay as shown in Fig. 5.7.

5.3 Practical Limits

In this section, we evaluate the ability of LISLs to provide reduced latency given practical implementation constraints. Specifically, we analyze the impact of two factors: the number of terminals available on the satellite as well as the impact of the processing delay.

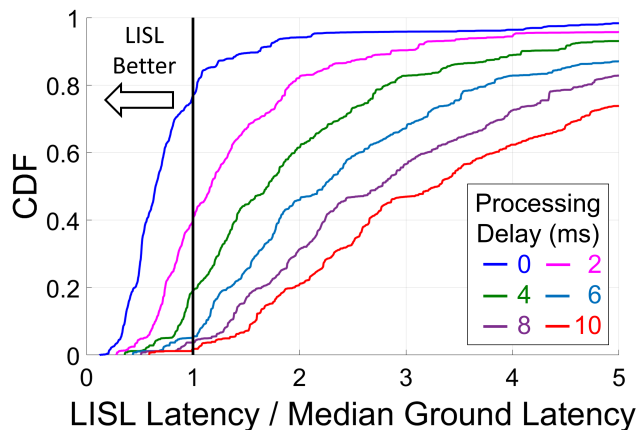


Figure 5.7: Practical Limits (a) Processing delay Impact: The efficacy of LISL-based constellations to provide latency benefits drops significantly with increasing values of processing delays.

Impact of processing delay on latency benefits: An important factor that will impact the behavior and design choices in the formation of the LISL constellations is the processing delay. While increasing the number of satellites can typically improve coverage, each layer/geometry plane adds an additional node whose processing and queuing delays are not addressed by the above propagation delay analysis. Thus, we evaluate the impact of this important system parameter on the benefit of LISL-based HFT transactions. The results in Fig. 5.7 show that as the processing delay on each node/satellite increases, the benefits of leveraging LISLs drop significantly. Even with small processing delays of 2 ms and 4 ms, the efficacy of providing benefit over terrestrial links falls from 78.76% to 40.5% and 19.05% respectively. In fact, a practical system with an on-satellite delay of 10 ms with the current topology becomes relatively ineffective in improving upon traditional terrestrial links.

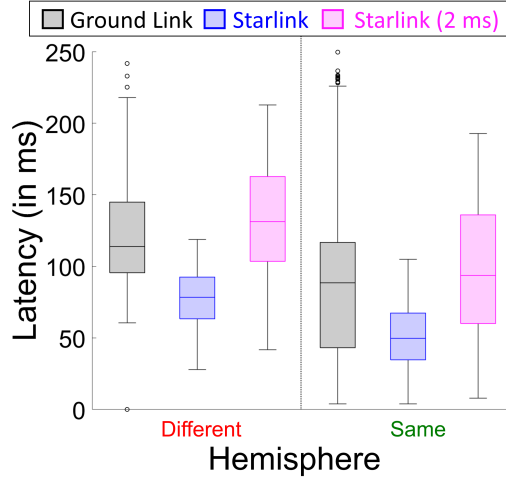


Figure 5.8: Practical Limits (b) Impact of geographical locations: Similar to terrestrial links, constellations face relatively higher median latency for connectivity across hemispheres which will be significantly increased with processing delay.

Impact of geographical locations: Further, the results in Fig. 5.8 show that, in the absence of processing delays, LISL-based constellations achieve significant benefits compared to terrestrial networks. However, when the processing delay is considered, the LISL-based latency will increase significantly due the increased use of IOP links to traverse more distance in the latitudinal direction. In fact, the evaluation with a processing delay of 2 ms shows how this significantly reduces the benefits achieved.




Figure 5.9: Practical Limits (c) Impact of terminals on latency benefit: In the presence of processing delay, NOP links show a definitive latency advantage over just IOP + AOP LISLs.

Impact of terminals on latency benefits: One surprising observation from the above evaluation is that almost all LISL routes utilize the IOP and AOP links primarily since NOP links are roughly the same latency as two AOP links (without processing delay). However, in the presence of a processing delay of 2 ms per satellite, this assumption no longer holds. Further, recent reports of deployed laser terminals demonstrate that only four terminals can be deployed on every satellite [45]. Thus, we endeavor to ascertain how much latency benefit is lost due to this constraint, i.e., how much latency benefit would two additional laser terminals provide in the presence of processing delays on the satellites. The result in Fig. 5.9 demonstrates that LISL median latency benefit can be increased by 15.4% while the number of link-pairs which receive latency benefits over terrestrial links can be improved by 24.6% from 40.5% to 50.47%.

Chapter 6

Discussion

In this section, the limitations of this empirical study are discussed and we describe some important open problems in LISL-based constellations for other researchers:

Limitations of Trace-driven Evaluation: Much of the evaluation in this study was limited to trace-based simulations due to the lack of publicly available technologically-mature testbeds. Perfect constellations were modeled, but are non-existent in deployment. Around 5-7% of satellites fail before their expected lifespan [55]. Moreover, with an increase in satellite launches and crowding in LEO, collision-avoidance trajectory maneuvers are exponentially increasing and result in orbital misalignment which would require re-establishment of LISLs [53]. Even the processing and queuing delays evaluated are chosen by scaling comparison to the fastest relays on terrestrial deployments. This may not be feasible in practical satellite based deployments. Further, aspects such as delay in LISL link formation and impact of satellite wobble on the latency have not been practically modeled. Finally, much of the discussions have ignored the quite significant MAC and upper layer latencies. This may further decrease the measured benefits of LISL-based networks compared to the terrestrial network latency.¹ Only permanent links are considered since establishing a new dynamic link can take between 10-30 seconds [56], which renders them incompatible with HFT.

¹Note, that despite this best case scenario for LISLs, they fail to improve over propagation delay in several demonstrated scenarios.

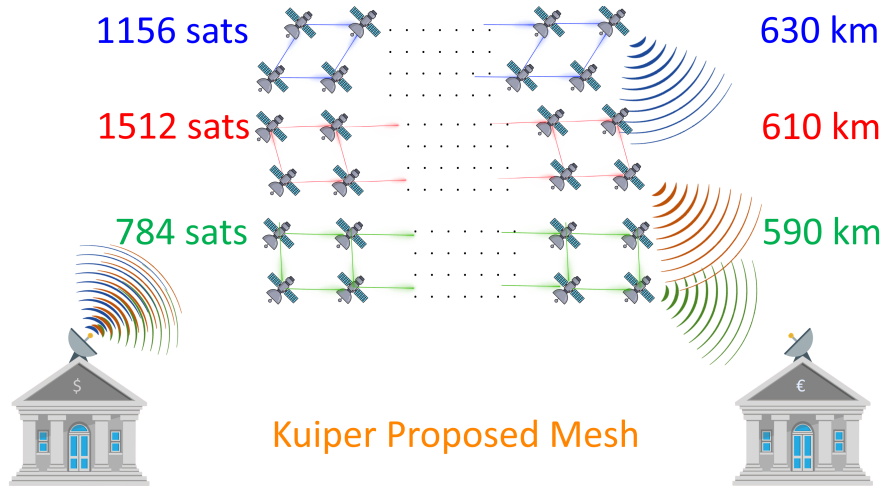


Figure 6.1: Kuiper is expected to deploy three satellite meshes with different altitudes and inclination angles in its constellation which will not be able to form permanent LISLs across them.

Multiple Mesh LISL networks: Both Kuiper (see Fig. 6.1) and Starlink have proposed to operate in a multi-mesh scenario which operates three to four separate meshes at different altitudes and inclination angles. However, given that each mesh operates in isolation, much of the analysis presented in this study does not translate to the practical scenario. The coverage provided by multiple meshes may be superior but the practical observed latency will remain within the bounds of empirical error presented. Note that as demonstrated in Sec. 5, there are several parameters such as propagation delay and phasing parameter that impact the coverage significantly more than the number of meshes surrounding Earth (except in potential congestion scenarios).

Deterministic Routing within the mesh: Given the known locations and traffic generated from these financial hubs as well as the rough location of the constellation across time, we can develop better predictive routing algorithms for the LISL mesh that can perform optimal latency/queuing delay trade-off. These algorithms can provide significantly better efficiency than terrestrial routing algorithms given the prior knowledge of temporal bandwidth requirements.

6.1 Take-aways

Based on the analysis, we describe the final recommendations to both the industry and research community:

Take-away 1 *New satellite constellation deployments should put emphasis on choosing the right values of the phasing parameter to provide better coverage.*

As new constellations get deployed across altitudes, they typically need to conform to the same constellation parameters as the incumbents to avoid collisions. Thus, it is quite critical for the initial incumbents to choose the right constellation parameters as they will significantly affect the coverage of future constellation deployments, and other competitors in the space industry in general. Further, this coverage benefit extends to providing traditional RF links and remains useful for incumbents.

Take-away 2 *LISL latency benefits are more pronounced between ground station pairs where the path through terrestrial infrastructure is longer than the orthodromic distance.*

The latency of the ground networking infrastructure across continents is quite significantly affected by the placement of the trans-oceanic cables to connect remote places. This certainly contributes to the high latency perceived via terrestrial links connecting the southern hemisphere. LISL meshes are unconstrained and equitable across geography, and hence provide significant latency benefits.

Take-away 3 *LISL network developers should prioritize reducing processing and queuing delays on satellites to below a millisecond to achieve significant latency benefits.*

The impact of the processing delay can be reduced by designing the topology to interpose AOP and NOP LISLs across a geometry plane of the constellation. Given the extensive prior knowledge about the LISL constellation mesh and amount of data generated from these financial hubs, intelligent and predictable routing mechanisms can achieve better efficiency than ground infrastructure.

Chapter 7

Conclusion

This work presented a holistic study on the feasibility and benefits of leveraging laser inter-satellite links (LISLs) for achieving low-latency high-frequency trading (HFT) transactions. This study takes an empirical approach in identifying the coverage of two of the largest deployed and proposed satellite constellations, Starlink and Kuiper, measuring the benefit of leveraging these constellations and the practical limits of these links. The evaluation across 157 million satellite-seconds of trace-driven simulation shows that LISLs can provide full coverage with roughly 1000 satellites and improve latency for 78.76% of link-pairs. However, when processing delays increase, the number of link-pairs showing improvement also drops significantly (e.g., at a processing delay of 4 ms per hop, only $\sim 19.05\%$ of the link-pairs benefit from this approach). Finally, we describe several open research problems that can overcome these practical limits by discovering better deterministic LISL-mesh routing algorithms, more complex LISL-focused topology design, and faster processing on the satellites.

Chapter 8

Future Work

To assess a best-case scenario for LEO constellation networks, transmission delays and link congestion for the LISL network were not modeled in this framework. Because the majority of HFT transactions use the FAST protocol [39] built on top of UDP and the link capacity per LISL terminal is purported to only be 10 Gbps [57], future work may include further modeling that accounts for this. Moreover, current HFT FPGAs have a very low tolerance for jitter, i.e., the statistical variance in latency. Modeling link congestion will invariably affect jitter and its impact should also be investigated. Inclusion of these factors will provide a fuller picture of the feasibility. Priority-queuing for HFT transactions to workaroud congestion and its limits should also be examined.

Distributed routing algorithms play a crucial role in realizing the benefits of an LEO constellation with LISLs. Future modeling could take this into account and pre-calculate path weights with all nodes in the system having coherent routing knowledge in a time-sensitive fashion. Precise timing for path-change and hand-offs known in advanced could minimize processing delay at each node and for the system as a whole.

Bibliography

- [1] T. Ahmmed, A. Alidadi, Z. Zhang, A. U. Chaudhry, and H. Yanikomeroglu, “The digital divide in Canada and the role of LEO satellites in bridging the gap,” *IEEE Communications Magazine*, vol. 60, no. 6, pp. 24–30, 2022.
- [2] A. Osipovich, “High-frequency traders eye satellites for ultimate speed boost,” *Wall Street Journal*, 2021.
- [3] “Starlink Business,” <https://www.starlink.com/business>, accessed: 2023-06-30.
- [4] M. Handley, “Delay is not an option: Low latency routing in space,” in *Proceedings of the 17th ACM Workshop on Hot Topics in Networks*, 2018, pp. 85–91.
- [5] A. U. Chaudhry and H. Yanikomeroglu, “Laser intersatellite links in a starlink constellation: A classification and analysis,” *IEEE Vehicular Technology Magazine*, 2021.
- [6] —, “Temporary laser inter-satellite links in free-space optical satellite networks,” *IEEE Open Journal of the Communications Society*, vol. 3, pp. 1413–1427, 2022.
- [7] D. Calzolaio, F. Curreli, J. Duncan, A. Moorhouse, G. Perez, and S. Voegt, “EDRS-C—The second node of the European Data Relay System is in orbit,” *Acta Astronautica*, vol. 177, pp. 537–544, 2020.
- [8] D. Bhattacharjee, A. U. Chaudhry, H. Yanikomeroglu, P. Hu, and G. Lamontagne, “Laser inter-satellite link setup delay: Quantification, impact, and tolerable value,” in *2023 IEEE Wireless Communications and Networking Conference (WCNC)*. IEEE, 2023, pp. 1–6.

- [9] S. Vatansever and I. Butun, “A broad overview of gps fundamentals: Now and future,” in *2017 IEEE 7th Annual Computing and Communication Workshop and Conference (CCWC)*. Ieee, 2017, pp. 1–6.
- [10] A. Al-Hourani, “A tractable approach for predicting pass duration in dense satellite networks,” *IEEE Communications Letters*, vol. 25, no. 8, pp. 2698–2702, 2021.
- [11] “Telesat Inflight Connectivity,” <https://www.telesat.com/inflight-connectivity/>.
- [12] “Application to Modify Petition for Declaratory Ruling to Grant Access to the U.S. Market for Telesat’s NGSO Constellation,” <https://fcc.report/IBFS/SAT-MPL-20200526-00053/2378318.pdf>.
- [13] “Global Communications | Services, Solutions & Satellite Internet | Viasat,” <https://www.viasat.com/>.
- [14] “FCC Grants SpaceX’s Satellite Broadband Modification Application,” https://docs.fcc.gov/public/attachments/FCC-21-48A1_Rcd.pdf, accessed: 2023-06-30.
- [15] “FCC Authorizes Kuiper Satellite Constellation,” https://docs.fcc.gov/public/attachments/FCC-20-102A1_Rcd.pdf, accessed: 2023-06-30.
- [16] Z. Qu, G. Zhang, H. Cao, and J. Xie, “LEO satellite constellation for Internet of Things,” *IEEE access*, vol. 5, pp. 18 391–18 401, 2017.
- [17] Z. Lai, W. Liu, Q. Wu, H. Li, J. Xu, and J. Wu, “SpaceRTC: Unleashing the low-latency potential of mega-constellations for real-time communications,” in *IEEE INFOCOM 2022*, 2022, pp. 1339–1348.
- [18] Y. Ulybyshev, “Satellite constellation design for complex coverage,” *Journal of Spacecraft and Rockets*, vol. 45, no. 4, pp. 843–849, 2008.
- [19] S. W. Paek, S. Kim, and O. de Weck, “Optimization of reconfigurable satellite constellations using simulated annealing and genetic algorithm,” *Sensors*, vol. 19, no. 4, p. 765, 2019.

- [20] R. Deng, B. Di, H. Zhang, L. Kuang, and L. Song, "Ultra-dense leo satellite constellations: How many leo satellites do we need?" *IEEE Transactions on wireless communications*, vol. 20, no. 8, pp. 4843–4857, 2021.
- [21] J. Liang, A. U. Chaudhry, and H. Yanikomeroglu, "Phasing parameter analysis for satellite collision avoidance in starlink and kuiper constellations," in *2021 IEEE 4th 5G world forum (5GWF)*. IEEE, 2021, pp. 493–498.
- [22] H. Hauschildt, S. Mezzasoma, H. L. Moeller, M. Witting, and J. Herrmann, "European data relay system goes global," in *IEEE ICSOS*, 2017.
- [23] S. Müncheberg *et al.*, "Development status and breadboard results of a laser communication terminal for large LEO constellations," in *SPIE ICSO*, 2019.
- [24] G. Zeng, Y. Zhan, H. Xie, and C. Jiang, "Networked satellite telemetry resource allocation for mega constellations," in *IEEE ICC*, 2022.
- [25] O. Markovitz and M. Segal, "Advanced routing algorithms for low orbit satellite constellations," in *ICC 2021-IEEE International Conference on Communications*. IEEE, 2021, pp. 1–6.
- [26] Z. Lai, H. Li, and J. Li, "Starperf: Characterizing network performance for emerging mega-constellations," in *IEEE ICNP*, 2020.
- [27] Q. Zhu, H. Tao, Y. Cao, and X. Li, "Laser inter-satellite link visibility and topology optimization for mega constellation," *Electronics*, vol. 11, no. 14, p. 2232, 2022.
- [28] K. An, M. Lin, T. Liang, J.-B. Wang, J. Wang, Y. Huang, and A. L. Swindlehurst, "Performance analysis of multi-antenna hybrid satellite-terrestrial relay networks in the presence of interference," *IEEE Transactions on Communications*, vol. 63, no. 11, pp. 4390–4404, 2015.
- [29] D. Bhattacharjee and A. Singla, "Network topology design at 27,000 km/hour," in *Proceedings of the 15th International Conference on Emerging Networking Experiments And Technologies*, 2019, pp. 341–354.

- [30] “Ham Radio Enthusiasts vs. High-Frequency Traders: A Battle for the Airwaves,” <https://www.wsj.com/articles/ham-radio-enthusiasts-vs-high-frequency-traders-a-battle-for-the-airwaves-6d6c5c1f?ns=prod/accounts-wsj>.
- [31] P. W. Lemme, S. M. Glenister, and A. W. Miller, “Iridium (r) aeronautical satellite communications,” *IEEE Aerospace and Electronic Systems Magazine*, vol. 14, no. 11, pp. 11–16, 1999.
- [32] S. Moorhead, “The iridium satellite network,” *Journal of Telecommunications and the Digital Economy*, vol. 10, no. 1, pp. 107–112, 2022.
- [33] Y. Zhang, Q. Wu, Z. Lai, and H. Li, “Enabling low-latency-capable satellite-ground topology for emerging LEO satellite networks,” in *IEEE INFOCOM 2022*, 2022, pp. 1329–1338.
- [34] C. M. Jones, “What do we know about high-frequency trading?” *Columbia Business School Research Paper*, no. 13-11, 2013.
- [35] R. Martin, “Wall street’s quest to process data at the speed of light,” *Information Week*, vol. 4, no. 21, p. 07, 2007.
- [36] C. Duhigg, “Stock traders find speed pays, in milliseconds,” *Sat*, vol. 1, no. 05, 2009.
- [37] M. Aitken, D. Cumming, and F. Zhan, “Trade size, high-frequency trading, and colocation around the world,” *The European Journal of Finance*, vol. 23, no. 7-9, pp. 781–801, 2017.
- [38] “The Secret World of Microwave Networks,” <https://arstechnica.com/information-technology/2016/11/private-microwave-networks-financial-hft/>.
- [39] C. Leber, B. Geib, and H. Litz, “High frequency trading acceleration using FPGAs,” in *2011 21st International Conference on Field Programmable Logic and Applications*. IEEE, 2011, pp. 317–322.
- [40] J. G. Walker, “Satellite Constellations,” *Journal of the British Interplanetary Society*, vol. 37, p. 559, Dec. 1984.
- [41] H. Kaushal and G. Kaddoum, “Optical communication in space: Challenges and mitigation techniques,” *IEEE communications surveys & tutorials*, vol. 19, no. 1, pp. 57–96, 2016.

- [42] A. U. Chaudhry and H. Yanikomeroglu, “Free space optics for next-generation satellite networks,” *IEEE Consumer Electronics Magazine*, vol. 10, no. 6, pp. 21–31, 2020.
- [43] V. W. Chan, “Free-space optical communications,” *Journal of Lightwave technology*, vol. 24, no. 12, pp. 4750–4762, 2006.
- [44] H. Kaushal, V. Jain, S. Kar, H. Kaushal, V. Jain, and S. Kar, “Acquisition, tracking, and pointing,” *Free Space Optical Communication*, pp. 119–137, 2017.
- [45] “SpaceX Non-Geostationary Satellite System Attachment A Technical Information To Supplement Schedule S,” <https://fcc.report/IBFS/SAT-MOD-20181108-00083/1569860.pdf>.
- [46] “Global Ping Statistics,” <https://wondernetwork.com/pings>, accessed: 2023-06-30.
- [47] D. J. Israel, B. L. Edwards, R. L. Butler, J. D. Moores, S. Piazzolla, N. du Toit, and L. Braatz, “Early results from NASA’s laser communications relay demonstration (LCRD) experiment program,” in *Free-Space Laser Communications XXXV*. SPIE, 2023.
- [48] M. H. Lane and F. R. Hoots, “General perturbations theories derived from the 1965 lane drag theory,” 1979.
- [49] R. S. Hujsak, “A restricted four body solution for resonating satellite with an oblate earth,” in *American Institute of Aeronautics and Astronautics Conference*, 1979.
- [50] “NORAD GP Element Sets,” <https://celestrak.org/NORAD/elements/gp.php?GROUP=starlink&FORMAT=tle>, accessed: 2023-06-14.
- [51] “Request for Modification of the Authorization for the SpaceX NGSO Satellite System,” <https://apps.fcc.gov/els/GetAtt.html?id=259301>.
- [52] E. W. Dijkstra, “A note on two problems in connexion with graphs,” *Numerische mathematik*, vol. 1, no. 1, pp. 269–271, 1959.
- [53] “SpaceX Starlink satellites had to make 25,000 collision-avoidance maneuvers in just 6 months — and it will only get worse,” <https://www.space.com/starlink-satellite-conjunction-increase-threatens-space-sustainability>.

- [54] A. U. Chaudhry and H. Yanikomeroglu, “When to crossover from earth to space for lower latency data communications?” *IEEE Transactions on Aerospace and Electronic Systems*, 2022.
- [55] “‘Not Good Enough’ - SpaceX Reveals That 5 Failed In Orbit So Far,” <https://www.forbes.com/sites/jonathanocallaghan/2019/06/30/not-good-enough-spacex-reveals-that-5-of-its-starlink-satellites-have-failed-in-orbit/?sh=656e18607e6b>.
- [56] C. Carrizo, M. Knappek, J. Horwath, D. D. Gonzalez, and P. Cornwell, “Optical inter-satellite link terminals for next generation satellite constellations,” in *Free-Space Laser Communications XXXII*. SPIE, 2020.
- [57] M. Motzigemba, H. Zech, and P. Biller, “Optical inter satellite links for broadband networks,” in *2019 9th International Conference on Recent Advances in Space Technologies (RAST)*. IEEE, 2019, pp. 509–512.