Scalability of LEO Satellite Constellations with Static and Dynamic Topologies

Robert Esswein, Quincy Bayer, Mai Abdelhakim, Robert K. Cunningham
University of Pittsburgh
Pittsburgh, PA USA
{robert.esswein, qub5, maia, robertkcunningham}@pitt.edu

***Abstract*—Low Earth Orbit (LEO) satellite constellations are becoming key enablers for connectivity and broadband coverage, specifically in underserved or hard-to-reach areas. LEO satellites are cheaper to launch and build, offer lower latency communication, and can be replaced more easily than Geostationary (GEO) satellites. However, LEO satellite systems face several challenges. LEO constellations require many more satellites to cover the earth, necessitating communication throughout the constellation. In addition, the coverage of each LEO satellite changes as it orbits the earth, unlike GEO, making constellation design and control more challenging. In this paper, we analyze the network performance of LEO satellite constellations. We focus on propagation delays, a major delay component in inter-satellite communications due to the vast distances. We examine the impact of the number of orbital planes and the number of satellites per orbital plane on the network performance. We analyze two different topologies, static and dynamic, and apply two different routing protocols for satellite communications. We show that, for a fixed number of satellites, additional satellites per orbital plane achieves superior performance compared to increasing the number of orbital planes in a static topology. However, for the dynamic topology, arranging the satellites in more orbital planes yields better performance. In any general, the dynamic topology has superior performance compared to the static topology. This paper sheds light on the impact of scalability and architecture of satellite orbits on the network performance of LEO satellite constellations, highlighting the benefits of the dynamic topology.**

I. Introduction

Satellite constellations are becoming more widely adopted for observation, navigation, and communication applications [1]. Compared to ground-based networks, satellite constellations have several key advantages. Once in orbit, satellites are able to service remote locations comparatively easily. Satellites are also self-contained systems, providing power generation and distribution, flight controls, communication, and application-specific functionality. This adoption of satellite constellations for numerous applications coincides with the trend of new constellations being designed for Low Earth Orbit (LEO) at an altitude of 500-2000 km instead of Geostationary Orbit (GEO) at an altitude of 36,000 km. LEO satellites have advantages when compared to GEO: communication latency is lower due to the lower altitude, and launching to orbit is less expensive; it is about half the price of GEO satellites [2]–[4]. LEO satellites also have drawbacks: LEO satellites do not orbit above a consistent point on earth's surface, which means that a satellite's ground coverage changes with time. LEO satellites also are unable to cover as large of a region on the earth's surface. For example, a GEO satellite can cover up to one third of earth's surface, while satellites in LEO orbit can typically only cover hundreds to a couple thousand km^2 [5], [6]. This means that multiple LEO satellites are required to service the same region as a GEO satellite and more advanced networking capabilities are required for the LEO satellite constellation.

This research was supported by the NSF Center for Space, High-performance, and Resilient Computing (SHREC) industry and agency members and by the IUCRC Program of the National Science Foundation under Grant No. CNS-1738783.

In recent years, we have seen several commercial satellite constellations; examples include Starlink [6] and Kuiper [7]. Both of these constellations will include over a thousand satellites. With this in mind, it is important to consider the implications of scaling the constellation. Satellites within one constellation can typically communicate with each other (inter-satellite communications) using laser/optical or RF communication channels. To facilitate communication between cities, data must travel through many satellites on route to the destination. This requires dedicated routing logic in all satellites. However, these constellations must be carefully designed for optimal network performance. As constellation designers, we must carefully decide between different constellation size parameters, routing algorithms, and topologies to achieve the required network performance.

In this paper, we analyze the network performance of satellite constellations and will examine how different characteristics of the satellite constellation network affect the communication latency. We focus on propagation delays since long distances between satellites significantly impact communication latency. We examine the impact of the constellation scale, specifically the number of orbital planes (OPs) and the number of satellites per OP, on the network performance. In addition, we examine both a static and a dynamic topology. We apply two different routing protocols for satellite communications to static and dynamic topologies. We show that for a fixed number of satellites, additional satellites per OP achieves superior performance compared to increasing the number of OPs in a static topology. However, for the dynamic topology arranging the satellites in more OPs rather than increasing the number of satellites per OP results in a better performance. This paper sheds light on the impact of the network architecture in satellite orbits on the network performance of LEO satellite constellations. We also

979-8-3315-2096-0/25/$31.00 ©2025 IEEE

show that the dynamic topology can achieve lower latency when compared to the static topology.

II. Related Work

Previous work examines how the delay of LEO satellite communication networks is affected by the altitude of the constellation [8] . This earlier work uses a satellite constellation model used in older systems, and only considers the smallest possible constellation for a particular altitude. Our work considers more satellites and the process of increasing the number of satellites in a constellation. This trend is useful for increasing bandwidth, decreasing latency, and is enabled by satellites become cheaper [3]. Additionally, we consider specific routing algorithms and topologies that were not available at the time of this publication.

Additional previous work in [9] examines large scale LEO satellite constellations with variable numbers of satellites. This work utilizes a downlink system model which assumes a Poisson point process distribution of satellites throughout the LEO region. Assuming a randomized distribution allows them to work with satellite density rather than any particular orbital pattern such as the common Walker-Delta constellations [10]. Additionally this work focuses on satellite to ground signals while we consider specifically satellite to satellite communications.

Previous work in [11] also explores the efficiency of communications in LEO satellite communications. However, this work analyzes communications between GEO and LEO satellites as a single server model rather than a LEO constellation. Additionally this work focuses on the antenna scheduling problem for individual satellites while we focus on the routing problem of an entire constellation.

III. Satellite Network Model

We model constellations as Walker-Delta constellations [10], of which Starlink [6] is a recent commercial adoption. The parameters for such constellations are $i : t/p/f$, where i is the inclination angle of the orbital planes (OPs), t is the total number of satellites, p is the number of OPs, and f is the relative spacing between satellites in neighboring planes. This implies a fifth parameter, the number of satellites per OP: t/p. All satellites within an OP are spaced evenly throughout the entire orbit. OPs are spaced evenly around the equator, resulting in good coverage of the space between slightly beyond $-i$ and i latitude, depending on the field of view of the satellites. The f parameter determines the phase offset between satellites in neighboring OPs. In a Walker-Delta constellation, OPs intersect with each other twice per orbit, which has potential for causing collisions between satellites. By changing the f parameter, the minimum distance between satellites as they pass by each other can be maximized, reducing the risk of collision [12]. Fig. 1 shows the orbital paths of satellites in a Walker-Delta constellation with eight OPs ($p = 8$) over an equirectangular projection of the earth.

We consider two different satellite network topologies: static and dynamic. In the static topology, each satellite maintains four permanent links, two to the nearest neighbors within its orbital plane and two to the satellites in the same phase of the adjacent orbital planes. This topology is consistent with existing work [13]. We define the dynamic topology, where each satellite maintains three permanent links and uses its fourth link to temporarily connect with passing satellites, similar to previous work [14], [15]. The first two permanent links connect a satellite to each of its neighbors within its orbital plane. The third permanent link connects to a satellite in a neighboring orbital plane, where satellites within one orbital plane alternate the direction of their links between east and west. So, these permanent links create a topology similar to the static topology, except half of the links between satellites in neighboring OPs are removed. These permanent links are used to ensure that the constellation remains connected.

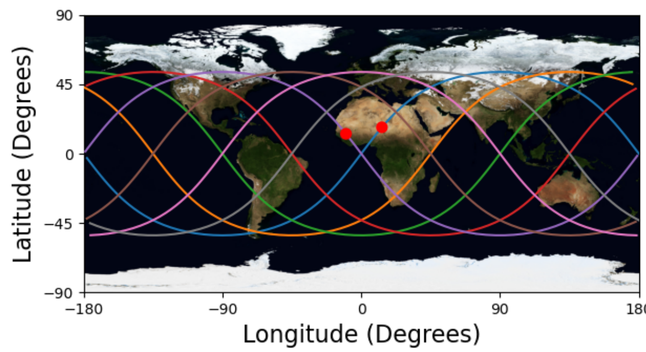

Fig. 1. The orbits of satellites on the equirectangular projection of the earth. Different colors help differentiate different OPs. Two dots are displayed showing how satellites might pass each other despite several OPs apart.

Satellites that are travelling in different directions near the same point would require data to travel very far for communication between them. For example, in Fig. 1, consider two satellites located at the points in the image. Using a static topology, communication between these two satellites would need to travel around the world even though they are physically close together. By allowing satellites to create temporary links, these two satellites might be able to connect directly, significantly reducing the distance communication must travel.

Since each satellite uses directional wireless optical links for communications, we must choose how satellites connect to each other. There are a variety of algorithms to select which temporary links are used [14], [16]. For this work, we use an algorithm that calculates how long different temporary connections are available, then greedily adds the longest connections into the constellation topology. This is a very simple method for selecting which temporary links to use, but we use this method to demonstrate the benefits of a dynamic topology. Some other algorithms might try to reduce the distance between cities [17] or create repeating patterns of connectivity [18].

IV. Constellation Routing

Routing algorithms developed specifically for satellite constellations exist, and two popular algorithms will be highlighted and examined in this paper. The first routing algorithm is the Orbit Prediction Shortest Path First (OPSPF) [19]. OPSPF builds on the Open Shortest Path First (OSPF) while utilizing the predictable nature of satellite constellations to improve efficiency. In OPSPF, each satellite periodically generates a routing table using a shortest path algorithm. This table is only valid for a short period of time, after which connections between satellites might go down or the shortest paths might change.

DisCoRoute [20] is another routing algorithm based on knowledge of the locations of satellites. However, DisCoRoute focuses on minimizing the computation required for routing. DisCoRoute assumes that the constellation uses a static topology; thus, this algorithm will only be tested for this topology. The key assumption of DisCoRoute is that hops between OPs tend to be shortest near the poles. This is because, looking at the slice of the earth at a particular latitude, the radius of this circle is smaller closer to the poles. DisCoRoute selects the path with inter-orbital plane hops are nearest to the poles, where links are shortest, approximating the shortest path.

V. Evaluation

In this work, we examine how the number of satellites in a constellation affects the network propagation delay. We focus on propagation delay because the vast distances travelled in inter-satellite communication makes it a major component of delay, while the high bandwidth links reduces the cost of other types of delay, such as transmission or queuing delay. There are two main parameters that we consider: the number of OPs and the number of satellites per OP. We examine constellations with different constellation parameters under both static and dynamic topologies, employing two different routing algorithms presented earlier in Section IV.

A. Experimental Setup

To evaluate the scalability of satellite constellation routing algorithms, we simulate these algorithms in the Mininet network simulator [21]. Both the OPSPF and DisCoRoute algorithms are simulated for numerous constellation sizes, and we measure the delay in communications between different cities. We select fifteen of the largest cities of high population density across the globe, with the additional constraint that all cities must be at least 20° degrees separated from any other city selected to ensure that inter-satellite communications are needed for the communications between any pair of cities [22]. Then traffic is generated from every city to every other city in this set, proportional to the sum of the populations of the two cities. The cities are shown in Fig. 2 and in Table I.

TABLE I
The cities used in this evaluation.

City Name	Country	Latitude (°)	Longitude (°)
Tokyo	Japan	35.68	139.69
Jakarta	Indonesia	-6.18	106.83
Guangzhou	China	23.13	113.26
Delhi	India	28.61	77.23
Mexico City	Mexico	19.43	-99.13
New York	USA	40.69	-73.92
Cairo	Egypt	30.04	31.24
Sao Paulo	Brazil	-23.55	-46.63
Lagos	Nigeria	6.46	3.38
London	England	51.51	-0.13
Los Angeles	USA	34.11	-118.41
Johannesburg	South Africa	-26.20	28.05
Dar es Salaam	Tanzania	-6.82	39.28
Lima	Peru	-12.06	-77.04
Moscow	Russia	55.76	37.62

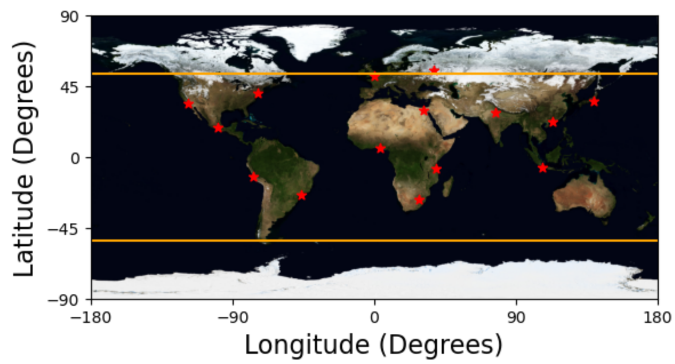

Fig. 2. Locations of cities used in evaluation on the equirectangular projection. Additionally, this graph shows lines marking the maximum and minimum latitude that satellites will fly over, though coverage is slightly beyond these lines because of the satellites' field-of-view.

In our simulation, each satellite has four laser transceivers to connect with nearby satellites. We assume that satellites have perfectly circular orbits, with no deviation from their ideal locations. In the dynamic topology, we assume that lasers are properly positioned after five seconds from ending the previous connection, and do not lose connection or interrupt communication until the links are brought down. We assume that communication between satellites is limited by the atmosphere and not by limitations of the lasers, which is consistent with several commercially available lasers [23]. We also assume that each satellite has a field of view of 60° towards the earth. We consider constellations at an altitude of 550 km, about a 90-minute orbital period, and a 53° inclination angle, all of which are consistent with Starlink constellation [6].

When looking at a satellite constellation, an important feature is the distance between satellites. Calculating the distance between two adjacent satellites within one OP is simple since the satellites are spaced evenly. So, two neighboring satellites are $360°p/t$ apart. The distance between connected satellites in adjacent OPs is more complicated. The distance between these satellites varies as the satellites orbit, though the satellites tend to be closest together near the poles. With some spherical trigonometry, we can show that on average the separation between connected satellites in neighboring OPs is less than $360°/p$.

A Walker-Delta constellation with sufficiently many satellites provides full coverage between $\pm(i+\delta)$ latitude, where i is the inclination angle of the orbital planes (OPs) and δ is determined by the field-of-view of the satellites. This is because the highest and lowest points of the orbits are at $\pm i$ and satellites are able to view a region on the ground. The highest and lowest latitude lines for the constellations we test, $\pm 53°$, are shown in Fig. 2. We notice that Moscow is slightly outside of this region, but it is still able to be covered by the constellation because it is always within a satellite's field of view. Reducing the total number of satellites in the constellation below a certain limit would result in some cities not being serviceable by a satellite at all times during the satellite orbits. This property was found by testing satellite constellations until one was found that failed to have a satellite in range of all cities at any time. This imposes a minimum number of 384 satellites in our experiments.

At any given moment in time, a point on the earth's surface might be within the field-of-view of multiple satellites. When a packet is generated at a particular location, we assume that the ground station is able to intelligently select the source and destination satellites. That is, a ground station can select the source-destination pair with the lowest cost to route traffic through. For the OPSPF algorithm, the satellites with the lowest distance between them is selected. Note that the DisCoRoute algorithm does not track the distance between nodes, instead the route with the fewest number of hops is selected. We do not consider the propagation delay for uplink or downlink, and focus on the inter-satellite routing.

Our simulator utilizes Mininet [21] to simulate network characteristics with updates to satellite locations every second. This leads to a small amount of error in the propagation delay calculation, however it is negligible compared to the total propagation delay of the system. At 550 km altitude, satellites are traveling at 8 km/s, so two satellites traveling in opposite directions sill separate at 16 km/s. This results in a maximum possible error in delay of 54 µs/s/hop. As we will see, the typical propagation delay is on the order of a few tens milliseconds, three orders of magnitude higher than the maximum possible error.

B. Impact of the Number of Orbital Planes

Our first experiment focuses on the impact of the number of orbital planes on the communications delay between cities through the constellation. We keep the number of satellites per OP constant at 24 for all constellations, then changed the number of OPs between 16, 18, 20, 22, 24, 26, and 28 OPs. This range was selected because all constellations provide full coverage of the earth and this range provides insight into how satellite constellations scale. The results are shown in Fig. 3.

It is observed that there is an optimal number of OPs (between 20 and 24 OPs) to achieve the lowest propagation delay between cities. One might expect that as we increase the number of satellites, more direct routes become available, allowing the constellation to reduce the propagation delay. However, that is not always the case because the propagation delay can only be reduced to a theoretical limit. Ultimately, cities are physically separated from each other, so propagation delay is limited by the speed of light and this distance. Approaching this theoretical minimum value is achieved by increasing the density of satellites, allowing for a shorter path. Beyond a certain density in the satellite constellation, there is no longer much benefit in adding more OPs because there is little room for improvement. The trends shown in this graphs can also be explained by looking at the number of hops and the per-hop distances which will be presented in Subsection V-D. The delay decrease as the number of hops decreases or the average per hop distance decreases.

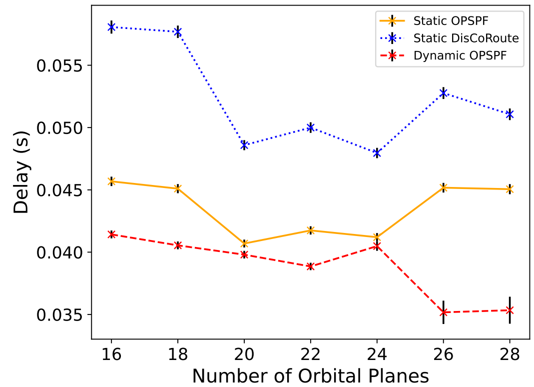

Fig. 3. The mean delay between cities for different constellation configurations and routing algorithms. The error bars show the 95% confidence interval for each point. This shows that the dynamic topology outperforms the static topology for all parameters.

From Fig. 3, comparing DisCoRoute to OPSPF on the static topology, we see very similar trends for all cases. As the number of OP increases, both of these algorithms see decreasing average propagation delay between cities. It can be seen that the OPSPF achieves a much lower delay compared to DisCoRoute for the different number of orbital planes. Additionally, as expected the dynamic topology outperforms the static topology. This is expected since the dynamic topology has one flexible link that can reduce the overall delay.

C. Impact of Number of Satellites Per Orbital Plane

In this experiment, we examine the impact of the number of satellites per OP on the delay performance of inter-satellite communications. Here, the number of OPs is kept constant at 24, and the number of satellites per OP changes between 16, 18, 20, 22, 24, 26, and 28 satellites per OP. Again, all of these constellations provide full coverage of the earth, which is why this range is used. The results from this test are shown in Fig. 4. As expected, increasing the total number of satellites in the constellation tends to decrease the propagation delay between cities. Similar to the previous experiment, we can see that OPSPF is superior to DisCoRoute, and dynamic topology helps reduce the communications latency.

It is noted that increasing the number of satellites per OP reduces the distance between satellites within one OP, while increasing the number of OPs reduces the distance between satellites in neighboring OPs. As discussed in Subsection V-A, the distance between satellites within one OP is $360°p/t$. Additionally, the maximum distance between satellites in neighboring OPs is slightly larger than $360°/p$, with the average distance being smaller. When the number of OPs is equal to the

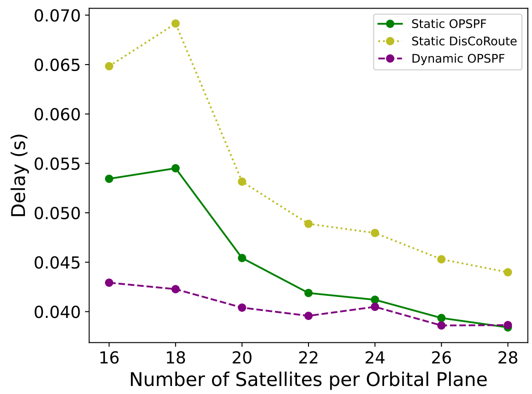

Fig. 4. The mean delay between cities for different constellation configurations and routing algorithms. The error bars show the 95% confidence interval for each point. This shows how increasing the number of satellites per OP reduces the propagation delay.

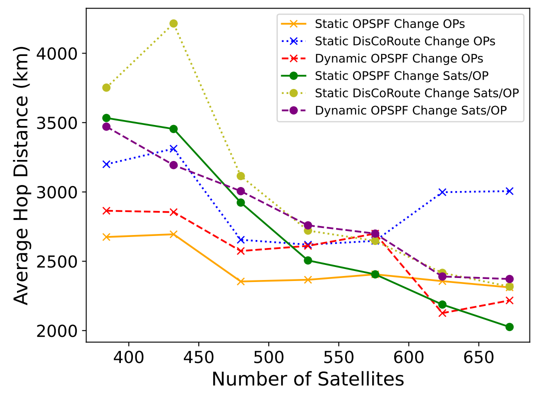

Fig. 6. The average distance traveled per hop. This shows that, even though the number of hops tends to increase, the distance traveled per hop tends to decrease as the size of the constellation increases.

number of satellites per OP ($p = t/p$), increasing the number of OPs reduces the distance between satellites in neighboring OPs. This dimension is already on average smaller than the distance between satellites within OPs.

D. Number of Hops

Fig. 5 shows how the number of hops is affected by the number of nodes in the network through increasing the number of OPs or the number of satellites per OP. In most cases, the average number of hops increases as the number of satellites increases. The region covered is constant, so signals transit more satellites between locations on the earth. In Fig. 6, we show the average distance per hop. In most cases, as the number of nodes increases, the distance traveled per hop decreases, except for DisCoRoute, where as the number of OP increases such that the total number of satellites exceeds 550, the number of hops decreases and the per-hop distance increases.

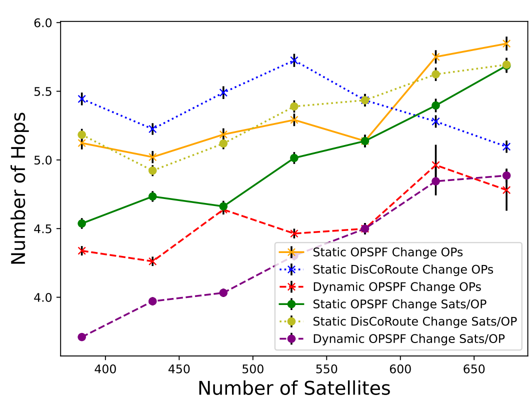

Fig. 5. The mean number of hops between locations on the ground. The error bars show the 95% confidence interval for each point. In most cases, as the number of satellites in the constellation increases, the number of hops between ground locations increases.

E. Combined Results

Lastly, we combine all results onto a single graph so that scaling the number of OPs can be compared with scaling the number of satellites per OP. Both experiments use the same total number of satellites in each constellation: 384, 432, 480, 528, 576, 624, and 672 satellites (which is why we use the same constant value of 24 in both experiments). The combined results are shown in Fig. 7.

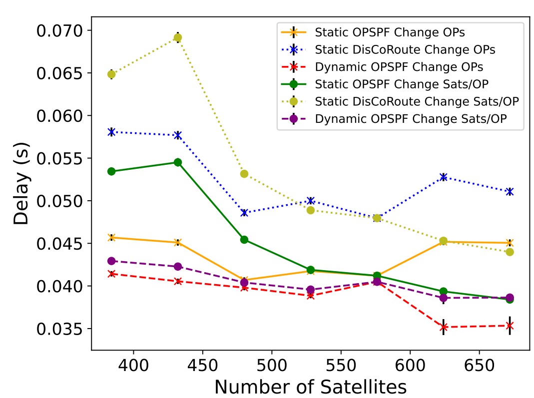

Fig. 7. The mean delay between cities for different constellation configurations and routing algorithms. The error bars show the 95% confidence interval for each point. This shows that the impact of changing the number of OPs and the number of satellites per orbital plane depends on the network topology.

As can be seen in Fig. 7, for the dynamic topology arranging the satellites in more OPs will result in a slightly better performance compared to increasing the number of satellites per OP, given a fixed total number of satellites in both cases. In the static topology, for a fixed number of satellites (greater than 600), additional satellites per orbital plane achieve superior performance compared to increasing the number of OP. When increasing the size of the constellation, the dynamic topology has the lowest latency for almost every test case since the dynamic topology allows for satellites to quickly send data in any direction. In the static topology, when satellites are traveling northeast, they have connections to the east, west, northeast, and southwest. When traveling southeast, satellites have connections to the east, west, southeast, and northwest, so it is costly for satellites to send data in a direction other than the direction of its links. For example, a satellite travelling southeast trying to send a packet northeast would first have to send the packet to the northwest then to the east. The dynamic topology gives each satellite up to one neighboring satellite that is traveling in the other direction. This means that even if a satellite is not connected to another satellite traveling in the opposite direction, it is likely near a satellite that has such a connection. So, the dynamic topology allows for satellites to more easily send data in any direction.

VI. Conclusion

In this paper, we analyze LEO satellite constellations communication latency. We assess the impact of the scale of satellite constellations on the propagation delay between different locations across the globe. We consider two different topologies, static and dynamic, and apply two different routing protocols for satellite communications. We show that for a fixed number of satellites, using more satellites per orbital plane achieves superior performance compared to a larger number of orbital planes in a static topology. This is different for dynamic topologies, where arranging the satellites in more orbital planes rather than increasing the number of satellites per orbital plane results in a better delay performance. We show that the dynamic topology provides a way to reduce the propagation delay across the network without requiring more satellites. Overall, the results show that the arrangement of satellites into orbits in space will depend on the network topology and the constellation routing algorithm used.

References

[1] Erik Kulu. “Satellite constellations-2021 industry survey and trends”. In: *Small Satellite Conference* (2021).

[2] Pedro L.K. da Cás, Carlos A.G. Veras, Olexiy Shynkarenko, et al. “A Brazilian space launch system for the small satellite market”. In: *Aerospace* 6.11 (2019), p. 123.

[3] Thomas G Roberts and Spencer Kaplan. “Space launch to low earth orbit: How much does it cost”. In: *Civil and Commercial Space Space Security* (2020).

[4] Stephen Clark. “Sizing up America’s Place in the Global Launch Industry”. In: *Spaceflight Now.* 24 (2013).

[5] Chenhua Sun, Huijuan Jiang, Jinsong Zhang, et al. “Modeling and Calibrating the Ground-Surface Beam Pointing of GEO Satellite”. In: *IEEE Access* 7 (2019), pp. 121897–121906.

[6] M Albulet. “Spacex non-geostationary satellite system: Attachment A technical information to supplement schedules”. In: *US Fed. Commun. Comm., Washington, DC, USA, Rep. SAT-OA-20161115-00118* (2019).

[7] FCC. *FCC authorizes Kuiper satellite constellation*. https://docs.fcc.gov/public/attachments/FCC-20-102A1.pdf. 2020.

[8] Bezalel Gavish and Joakim Kalvenes. “The impact of satellite altitude on the performance of LEOS based communication systems”. In: *Wireless Networks* 4 (1998), pp. 119–213.

[9] Zhaohua Qiu and Wen Wang. “Performance Analysis and Simulation of Large-Scale LEO Constellation Under a Stochastic Geometric Perspective”. In: *ICC 2023-IEEE International Conference on Communications* (2023).

[10] John G Walker. “Satellite constellations”. In: *Journal of the British Interplanetary Society* 37 (1984), p. 559.

[11] Y. Zhu et al. “Antenna Scheduling for Multiple User Satellites in Space Data Relay Networks”. In: *ICC 2019* (2019), pp. 1–7.

[12] Mark Handley. “Delay is not an option: Low latency routing in space”. In: *17th ACM HotNets*. 2018, pp. 85–91.

[13] Hui Li, Dongcong Shi, Weizheng Wang, et al. “Secure routing for LEO satellite network survivability”. In: *Computer Networks* 211 (2022), p. 109011.

[14] Qinyu Zhu, Xinmin Wang, Haitao Yang, et al. “Design and Evaluation of Dynamic Topology for Mega Constellation Networks”. In: *Electronics* 12.8 (2023), p. 1784.

[15] Robert Esswein, Sean O’Melia, Richard Skowyra, et al. “Poster: Resiliency in Low Earth Orbit Satellite Routing”. In: *IEEE Security & Privacy* (2023).

[16] Long Chen, Feilong Tang, Linghe Kong, et al. “Load-Adaptive and Energy-Efficient Topology Control in LEO Mega-Constellation Networks”. In: *GLOBECOM 2022*. 2022, pp. 1552–1557.

[17] Zhibo Yan, Guoyong Gu, Kanglian Zhao, et al. “Integer linear programming based topology design for GNSSs with inter-satellite links”. In: *IEEE Wireless Communications Letters* 10.2 (2020), pp. 286–290.

[18] Debopam Bhattacherjee and Ankit 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.

[19] Tian Pan, Tao Huang, Xingchen Li, et al. “OPSPF: orbit prediction shortest path first routing for resilient LEO satellite networks”. In: *International Conference on Communications (ICC)*. IEEE. 2019, pp. 1–6.

[20] Gregory Stock, Juan A Fraire, and Holger Hermanns. “Distributed On-Demand Routing for LEO Mega-Constellations: A Starlink Case Study”. In: *ASMS/SPSC*. IEEE. 2022, pp. 1–8.

[21] Bob Lantz, Brandon Heller, and Nick McKeown. “A Network in a Laptop: Rapid Prototyping for Software-Defined Networks”. In: *Proceedings of the 9th ACM SIGCOMM Workshop on Hot Topics in Networks*. 2010.

[22] Simple Maps. *World Cities Database*. https://simplemaps.com/data/world-cities. 2023.

[23] Tesat-Spacecom. *Laser Products*. https://www.tesat.de/products#laser. 2021.