A Call for Decentralized Satellite Networks

Seoyul Oh
University of Illinois Urbana-Champaign

Deepak Vasisht
University of Illinois Urbana-Champaign

ABSTRACT

Low Earth Orbit (LEO) satellite constellations are emerging as key to robust global internet connectivity, especially in areas that lack adequate terrestrial connectivity either due to lack of financial viability or due to disruptions caused by wars and natural disasters. Yet, such LEO constellations are few in number and can arbitrarily turn off access in times of conflict. This has led to demands for independent satellite constellations by different countries and organizations. We argue that such independent constellations are impractical, wasteful, and unsustainable due to the orbital dynamics of LEO satellites. Instead, we propose multi-party decentralized constellations wherein different parties contribute a small number of satellites to a shared constellation. Such multi-party constellations are robust to a subset of participants backing out and reduce economic costs, capacity waste, and orbital occupancy. We discuss multiple technical developments that make such designs possible today and list open questions for further investigation.

CCS CONCEPTS

• **Networks** → *Network design principles*; • **Computer systems organization** → *Dependable and fault-tolerant systems and networks.*

KEYWORDS

Satellite Networking, Decentralized Networks, Multi-party Constellations

ACM Reference Format:
Seoyul Oh and Deepak Vasisht. 2024. A Call for Decentralized Satellite Networks. In *The 23rd ACM Workshop on Hot Topics in Networks (HOTNETS '24), November 18–19, 2024, Irvine, CA, USA.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3696348.3696896

HOTNETS '24, November 18–19, 2024, Irvine, CA, USA

ACM ISBN 979-8-4007-1272-2/24/11
https://doi.org/10.1145/3696348.3696896

(a)

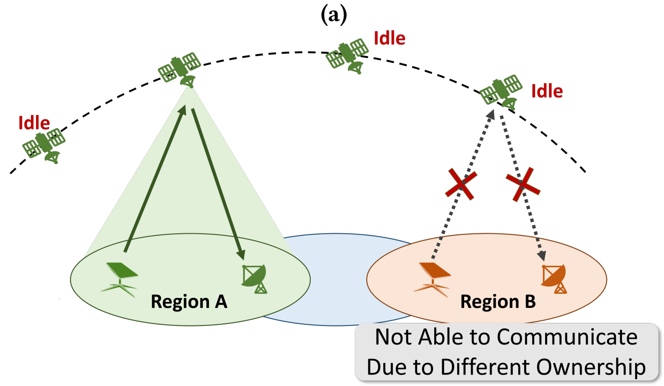

(b)

Figure 1: *(a) Orbital motion of a LEO satellite across three hours (the color changes from red to blue with time). (b) Designing region-specific satellite networks is inefficient. Satellites meant for region A are only available for few minutes in region A and have spare capacity during other times which can be used by other regions.*

1 INTRODUCTION

Low Earth Orbit (LEO) satellite constellations are experiencing rapid growth as means for global Internet connectivity. Such constellations are the primary means of connectivity in areas where terrestrial infrastructure is absent, insufficient, or destroyed by human or natural disasters (e.g., wars and natural emergencies). Recently, such connectivity has proven useful in many scenarios such as the tsunami in Tonga [35] and Russia's invasion of Ukraine [7].

Today, Starlink [37] is the largest deployed broadband constellation comprising nearly 6000 satellites [43], with over three million subscribers [49]. The primary competitors include OneWeb [32] and Amazon Kuiper [18] constellations. However, the centralized control of a few mega-constellations for such an essential utility is increasingly being seen with suspicion. Such suspicions were heightened by the tactical role of Starlink in the Russia-Ukraine war [6, 33, 36]. In response to such suspicions, Taiwan recently announced plans to build their own satellite network. Such steps are largely motivated by lack of trust in the operators – if the operator shuts down connectivity over a region during an emergency or war, users are left with little recourse. As a Taiwan government official notes [42], 'we cannot put all our eggs in one basket'. Similarly, South Korea announced plans to develop an operational LEO constellation by 2030 [41].

We argue that the launch of independent constellations by multiple nation states and corporations is (a) expensive, (b) wasteful, and (c) unsustainable. This is due to the unique orbital characteristics of LEO constellations. LEO satellites operate in low orbits, nearly 500-700 Kilometers above Earth, and have an orbital period of approximately 1.5 hours. Since Earth rotates at a period of 24 hours, these satellites are not in sync with Earth's motion. Therefore, as shown in Fig. 1a, the satellite covers different paths on Earth during each orbit. Due to this orbital motion, a single satellite can only offer few (less than ten) minutes of coverage per day to a given region. Providing complete coverage needs megaconstellations of hundreds to thousands of satellites, requiring expenditures of billions of dollars. Amazon and Starlink have projected that building fully operational LEO networks requires investments between 10-30 billion dollars [2, 34]. Despite such large investments, region-specific satellite networks waste most of their capacity because the satellite is doomed to be idle when it is not over the region of interest (e.g., in Fig. 1b, the satellites are idle when not over region A). We show in Sec. 2 that such satellites are idle over 99% of the time. Finally, an increase in the deployment of large constellations will lead to increased orbital congestion, with higher risks of collisions and increased obstructions for astronomical observations [44].

Therefore, we aim to identify an alternative solution for trusted and robust coverage without relying on independent satellite networks by each competing entity. Specifically, we propose a decentralized multi-party low earth orbit (MP-LEO) network. In an MP-LEO constellation, each participant contributes a small number of satellites and does not need to deploy large constellations to get continuous coverage. These satellites, however, offer their spare capacity to other users of the network when not in use by the contributor's devices (leading to reduced waste). In turn, the contributors can access the network through the spare capacity of other satellites. MP-LEO's goal is to design constellations that offer distributed control. Unlike the status quo, it is impossible for a single party to shut down the entire MP-LEO constellation or deny service to a specific region. Even if a small number of parties collaborate, they can cause (at worst) minor degradation to the network, proportional to their stake in the network. Therefore, MP-LEO can reduce economic costs and waste while improving trust and sustainability.

We envision that both government entities and private companies such as terrestrial internet service providers can participate in MP-LEO. Countries can guarantee satellite access during adversity through participation in such networks. Similarly, using MP-LEO, private companies can begin to serve as independent satellite network providers, without requiring multi-billion dollar investments, thereby increasing competition in this sector. Such multi-party collaborations have historically been successful in space, e.g., in the International Space Station [31]. Finally, there is a slew of emerging technological primitives that make such partnerships mutually beneficial and robust. For example, there is a recent emergence of satellite-as-a-service launch models where participants can just own or rent parts of a satellite (e.g., just the radio module for connectivity) to reduce launch costs [9, 40]. Similarly, advances in multi-party computation can enable robust distributed control [19, 22].

Designing MP-LEO opens up several dimensions of research questions. First, there are questions of designing incentives for participation. What financial incentives motivate new participants to join such networks? What constitutes *good behavior* for participating parties in such a shared network? How do satellite operators charge for their services? What are the *optimal* orbits for satellite deployment? The second dimension of questions revolve around trust and robustness. How do we prevent individual satellite operators from denying service to others while continuing to benefit from other satellites? What happens when a participating party chooses to step out of the network? How do we deal with satellite failures?

A key observation in our work is that incentive design and robustness of the MP-LEO constellations are closely related. For satellite networks, simply incentivizing individual parties to maximize coverage leads to more robust networks. Specifically, we find global coverage is optimized if we deploy satellites *far* from each other in distance (and orbital parameters). Such a design naturally leads to a constellation where satellites from multiple parties do not form a cluster and are interspersed. Such constellation designs are also more robust. Even if one party backs out from the constellation, it doesn't create large continuous gaps in orbital coverage and connectivity (but leads to reduced capacity).

We discuss these design considerations, trade-offs, and design choices in the rest of this paper. We note that the designs presented in this paper are works in progress, and highlight several open questions in Sec. 4.

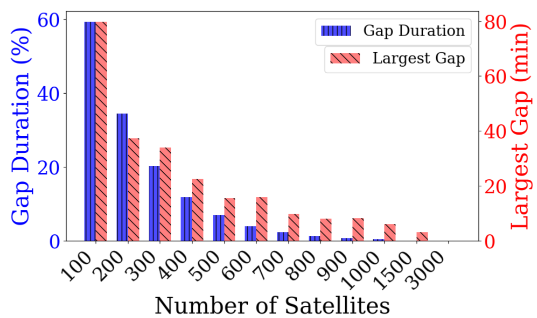

Figure 2: *Relationship between the number of satellites and percentage of time without coverage.*

2 THE NEED FOR DECENTRALIZED SATELLITE NETWORKS

We begin by quantifying the need for a decentralized constellation using Microsoft's CosmicBeats [29] simulator. It simulates satellite orbits and satellite-ground links using Two Line Element (TLE) orbit descriptors.

High cost of satellite networks: First, we ask how many satellites are needed to provide coverage to a target location, i.e., how many satellites would a country need to deploy to serve their own users? We select Taiwan as the target for this simulation and place a receiver at a central location in Taipei, Taiwan. We quantify the coverage gap across one week, averaged across one hundred runs of the simulation. In each run, we randomly sample satellites from the Starlink network.

Fig. 2 shows the percentage of time without any coverage (i.e., no satellite is visible) with increasing number of satellites. As shown, the coverage improves with increasing number of satellites in the constellation. With 100 satellites, the user has no coverage for over 50% of the time, with continuous gaps of up to over an hour. To achieve over 99.5% coverage, we need at least 1000 satellites. In practice, networks aim for five-nine (99.999%) availability, which would require even larger constellations.

Under-utilization due to orbital patterns: Next, we demonstrate that if a single region were to deploy their own satellite constellation, these satellites would be significantly underutilized outside their designated region. We use the simulation to quantify each satellite's idle time, i.e., times when it is not connected to a user terminal. For this simulation, we consider placing user terminals between one to 21 cities. The analysis includes the top 20 most populated cities, limited to one per country. We add Melbourne, Australia, to ensure representation from all major continents. Fig. 3 illustrates the relationship between the number of cities served and satellite idle time. If we deploy a constellation to serve just one major city, each satellite will be idle for 99% of the time, leading

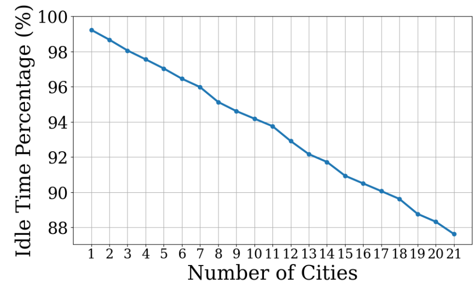

Figure 3: *Relationship between the number of cities served and satellite idle time*

to massive under-utilization. As we expand the number of cities covered across the world, satellite idle time decreases.

Our simulations confirm two key observations: (a) large satellite constellations and correspondingly large investments are necessary for providing coverage even in geographically small areas; (b) global sharing of satellites improves satellite utilization. This analysis also presents the potential upside for MP-LEO constellations – a participant contributing just 50 satellites can get coverage worth over 1000 satellites by trading off their spare capacities with others. These observations arise from the natural orbital motion of LEO satellites and highlight the need for decentralized networks. At this point, one might wonder – why not use geostationary satellites that do not move with respect to earth? Such satellites operate at heights of around 36000 Km, leading to orders of magnitude degradation in network latency (second-level) and capacity compared to LEO satellites.

Comparison to decentralized terrestrial networks: Such decentralized networks have been previously deployed in terrestrial cellular and Internet-of-things contexts, e.g., the Helium network [14]. We are motivated by these designs and believe that MP-LEO networks will build on many of the primitives in decentralized cellular networks. However, the orbital motion of satellites lead to a fundamental difference – a user cannot simply deploy satellites overhead to provide coverage to themselves. This is unlike terrestrial networks, where gaps in coverage can be filled with local solutions, e.g., by deploying additional base stations. Moreover, unlike the terrestrial case, the natural motion of satellites leads to long idle times when the satellite is not overhead. This makes decentralized satellite constellations more compelling than decentralized terrestrial networks.

Satellite design and launch: The design of satellites and their launches have increasingly been democratized, reducing the cost and expertise barriers of entry. This would allow even smaller companies and Internet service providers to participate in an MP-LEO constellation. For example, there is an emerging paradigm in the satellite community – satellite-as-a-service (SaaS) [9, 40]. Under this paradigm, a network operator can just design a radio module without having to worry about designing a satellite themselves. A satellite launch company will pack different modules (e.g., a camera module from Earth observation companies and a radio module from a MP-LEO participant) onto a single satellite to reduce launch costs. In addition, some companies [40] can rent out existing radio modules on already launched satellites. Such technological trends in the space sector enhance the feasibility of MP-LEO network designs.

3 CONSTELLATION DESIGN

Our design choices for MP-LEO are motivated by the following goals:

- **Global coverage:** To begin with, we aim to optimize satellite coverage for people across the globe.
- **Robustness and trust:** The constellation should be robust to individual parties leaving the network and satellite failures. Any degradation should be proportional to their stake in the network. Such robustness is essential for establishing multi-party trust.
- **Financial viability:** Network participants should be appropriately compensated for launching and maintaining satellites.
- **Incremental deployment:** We expect the network to be incrementally deployed, therefore, we consider that as the primary deployment case in our current analysis.

3.1 Design overview

A satellite network consists of three components: user terminals, satellites, and ground stations. Most satellite networks today operate using a bent-pipe architecture [8, 47]. A user terminal communicates data to a visible satellite, which relays it to a ground station in its range (see Fig. 1b). Some satellites today utilize inter-satellite links (ISLs) but such ISLs are infrequently used, e.g., when a ground station isn't visible.

For decentralized networks, we opt for a transparent bent-pipe architecture. In a transparent architecture, the satellite does not even decode the signal transmitted by the user terminal. The satellite acts as a radio-frequency repeater and simply repeats the uplink radio signal on its downlink to the ground station. This allows a satellite to support multiple communication protocols and modulations. Such transparent designs were commonly used in geostationary satellites because such satellites are in orbit for decades and cannot be updated to support new protocols frequently.

This design choice has three other benefits. First, terminals and ground stations can agree on their own local RF protocols, modulation schemes, and encryption, preventing satellites from decoding the signals and offering high privacy protections. Second, the burden of decoding, processing, and routing data falls on the ground stations, greatly simplifying the satellite design. Finally, this design does not require satellites to have ISL capabilities which would otherwise require interoperability among satellites from different operators.

In our design, a participant's terminals connect to their own ground stations. Given the recent emergence of ground-station-as-a-service offerings by major cloud providers like Amazon and Microsoft [1, 30], deploying a ground station does not require extensive effort and can be purely done by defining software-defined radio blocks that can run on existing ground station deployments. Therefore, if an internet service provider wants to offer services in a local area (say New York), they just need to deploy their terminals and ground stations (or rent existing ground stations in the area). On these sub-networks, the entities have complete control over billing, policy, and network security.

Finally, we note that such designs naturally circumvent concerns about data sovereignty due to the participants' ability to choose their own communication protocols and encryption schemes.

3.2 Participation Incentives

Terrestrial decentralized networks, such as Helium, provide financial incentives to network participants for (a) verifying the coverage being provided by other participants, (b) providing coverage to their users, and (c) verifying transactions in the network. In principle, MP-LEO networks can follow similar incentive structures. Ground stations at random locations can verify coverage by pinging satellites when they are overhead, and provide proof-of-coverage to earn rewards.

Similarly, consumers pay satellite operators to carry traffic, in proportion to utilization. These prices can be dynamically set, leading to open data markets, or they can be predetermined. Participants with more satellites have more opportunities to carry data and can, therefore, earn more money. These financial exchanges can be mediated by centralized or decentralized systems (e.g., cryptographic tokens). Note that, the same participant can both be a consumer (e.g., if they are using spare capacity of other satellites) and a provider (if their satellite's spare capacity is being used by others). The design of these markets, their equilibrium, and their bootstrapping are open questions that we discuss in Sec. 4.

One might argue that terrestrial p2p networks, such as Guifi [4], have succeeded without dependence on profitability or crypto-based incentives, demonstrating the potential for non-profit participatory models. However, without financial incentives, the fixed costs for setting up and maintaining a satellite node are significantly higher than those for terrestrial networks. This suggests that while participation incentives in satellite constellations can be inspired by these principles, they must also account for the steeper financial barriers inherent to satellite networks.

Helium-like networks design incentive structures to offer higher rewards in regions of low coverage. However, as discussed before, increasing coverage in specific areas is non-trivial in satellite networks. How does one design satellite constellations that can reduce coverage holes in space-time? We consider the implications of different orbits on such coverage holes in Sec. 3.3. Specifically, for the rest of the paper, we consider design decisions that maximize the population weighted coverage over 21 most populous cities in the world (picked as before – no two cities from the same country and the addition of Melbourne for coverage of the Australian continent). In practice, the constellations can optimize for other similar objectives, such as by accounting for purchasing power parity, cost of data, etc.

Finally, we note that participants in MP-LEO constellations can either choose to optimize for their profit (e.g., private companies) or optimize for connectivity in their own region (e.g., countries). In our simulations, we find that these choices are often co-related, but do not exactly lead to the same outcomes. Even when a participant optimizes for local gains over global outcomes, the spare capacity is spread across the globe and benefits the rest of the network.

3.3 Constellation Design to Reduce Coverage Gaps

Orbits are defined by their inclination with respect to the poles, height, and the orbital plane, and are constrained by physical laws as well as regulation. Given these constraints, we need to identify optimal satellite deployment such that it maximally increases the global coverage time. Simply adding a satellite anywhere in the constellation is insufficient to achieve this. We quantify the improvement in population-weighted global coverage time across one week, averaged over one hundred runs of the simulation. In each run, we randomly sample a single satellite from the Starlink network, and add it to the existing base satellite(s). We plot the results in Fig. 4a, which demonstrates the significant impact of satellite placement on coverage.

When adding a satellite to a single-satellite constellation, the coverage increases by over 1 hour on average, with a maximum increase over 4 hours. This enhancement is notable in lager constellations comprising 100 and 500 satellites. The gains in coverage are most pronounced when the initial number of satellites is small, highlighting the importance of strategic satellite positioning within the orbit to maximize coverage benefits.

To further investigate the optimal strategy for efficiently filling coverage gaps, we simulated an imaginary constellation with 12 satellites, each 30 degrees apart in the same orbital plane. This plane has an inclination of 53 degrees and a height of 546 Km (chosen to be consistent with Starlink's design). We then consider adding a satellite at 29 locations between two of these original satellites, spaced about 1 degree (120 km) apart in phase. We calculated the coverage improvement compared to the original set of 12 satellites. Adding a satellite at the midpoint between two original satellites (15 degrees from each) yielded the maximum coverage improvement, as shown in Fig. 4b. This demonstrates that strategically positioning a satellite at the farthest point from existing satellites maximizes coverage benefits.

Real-world constellation design also considers two other factors – inclination and altitude. To consider these factors, Fig. 4c shows the impact of adding a new satellite to an existing set of four Starlink satellites (53-degree inclination, spaced approximately 90 degrees apart in the same orbital plane). The new satellite was chosen from three categories: 1) Different inclination (43 degrees), 2) Same orbital plane and phase but different altitude, 3) Same orbital plane but different phase. The results showed that adding a satellite with a different inclination provided the highest coverage improvement, increasing by about 1 hour and 11 minutes. This suggests that varying inclination significantly improves coverage by introducing diverse orbital paths. Additionally, the other two categories (different altitude and phase) also showed over 30 minutes of improvement, suggesting that multiple factors can contribute to increased coverage.

Our analysis demonstrates that as the constellation evolves, contributors are incentivized to deploy satellites in orbital configurations different from other satellites. By identifying the largest coverage gaps and filling them, satellite operators will have exclusive access to potential consumers during these gaps increasing their revenues. Such placement also benefits the global network operation by maximizing network coverage. Finally, we note that this design also enhances robustness – when participants pull out, it does not lead to large continuous disruptions in coverage.

3.4 Robustness

How much does an MP-LEO constellation suffer in terms of coverage when participants decide to back out? We measured the coverage reduction over one week, averaging results from 100 simulation runs. In each run, we start with a base group of N satellites (the values of N are 200, 500, 1000, and 2000). Then, we withdraw a random selection of $\frac{N}{2}$ satellites. We estimate and plot the reduction in coverage in Fig. 5. When the satellite constellation is small ($N = 200$), the coverage reduction is significant at 1 day and 16 hours (24.17% drop in coverage). As the total number of satellites increases, this loss subsides – reducing to 0.37% at 2000 satellites. As expected, the network grows more robust as more satellites join in.

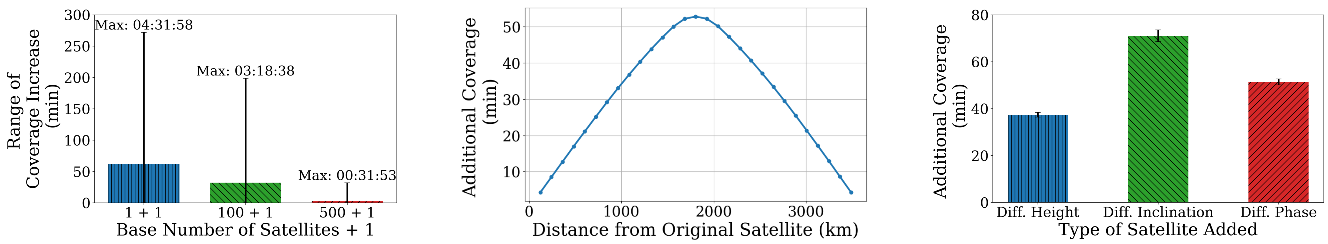

(a) *Addition of One Satellite* **(b)** *Impact of phase variation* **(c)** *Impact of different factors*

Figure 4: *Coverage Implications of Constellation Design – (a) A single satellite added to the constellation has decreasing benefits on the coverage provided by the constellation. (b) When a satellite is added between uniformly placed satellites, its impact is maximum when it is farthest from existing satellites. (c) Inclination, height, and phase have varying impacts on the additional coverage.*

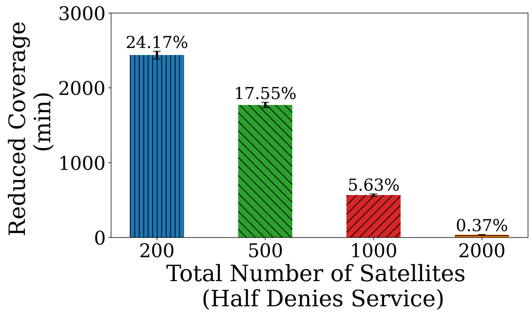

Figure 5: *Reduction in coverage when half of the satellites in a constellations deny service.*

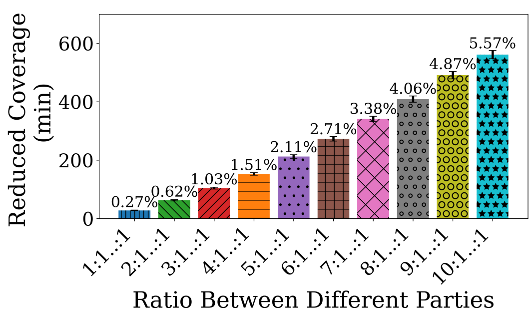

Figure 6: *Reduction in coverage (in percentage) for various ratios between different parties in MP-LEO networks. The largest party denies service.*

What is the impact on coverage when a single largest participant pulls out of an MP-LEO constellation? We simulated a scenario with 1,000 satellites where the contribution ratios of 11 parties vary from equal (1:1:...:1) to highly skewed (10:1:...:1). In each run, we withdrew the satellites of the largest party. The results in Fig. 6 show a clear correlation between contribution distribution and network robustness. When contributions are equal (e.g., 91 satellites each) the reduction in coverage is minimized, showcasing optimal resilience. Conversely, when contributions are skewed (e.g., one party contributes 500 satellites, others 50 each), the coverage loss is more pronounced. However, even in a highly skewed scenario, the network is service-able (5.5% gap or 10 hours of no coverage across a week) due to the other parties still contributing about half of the network.

4 OPEN QUESTIONS

Our goal is to leave the reader with multiple open questions about various aspects of MP-LEO design.

Bootstrapping decentralized networks: Bootstrapping is a challenging problem in decentralized networks. Early participants contribute a small number of satellites, which do not provide continuous coverage and, hence, find few customers. Such questions have been tackled by terrestrial decentralized networks by issuing tokens to early adopters with future financial value. In addition to exploration of token-based designs, we anticipate that early sparse MP-LEO deployments can provide global coverage for delay tolerant applications (e.g., IoT and opportunistic high volume transfers) at lower unit costs.

Market design: How much should satellite operators charge for data access? What kinds of quality-of-service can they provide? How do users choose between competing satellites after the deployment reaches complete coverage? These game theoretic explorations of market design are interesting open questions.

Bent-pipe architectures and inter-satellite links (ISLs): Different variants of MP-LEO's design can choose to alter the bent-pipe architectures to operate at the packet-level, i.e., the satellite decodes and relays bits as opposed to relaying the raw RF signal, thereby avoiding any amplification of noise from ground transmissions. Such designs can still offer packet-level encryption. Furthermore, our current design omits ISLs to simplify satellite architecture and reduce costs. However, future work can consider ISLs to enable data routing between satellites without needing to relay signals through ground stations.

Multi-party control: So far, we have not explored multi-party control, i.e., in our proposed design, each party controls their own satellites. Space-based trusted execution environments [28] have been proposed and can potentially be utilized to provide cryptographic guarantees on what runs on the satellite and how they are controlled (e.g., by consensus from multiple parties).

Open-source designs: Open-source designs for user terminals, satellite radios, and ground station receivers can greatly facilitate mass adoption. Such designs must account for size, weight, power, and RF constraints.

Spectrum access: Spectrum sharing between terrestrial and satellite users is an active topic of research, e.g., in 6G networks [17, 20]. Our current design delegates spectrum management to ground stations and user terminals since the satellite acts merely as a repeater (and will be designed as compatible with primary satellite frequencies – X and Ka/Ku bands). However, we anticipate that active spectrum management strategies are required for efficient spectrum utilization.

5 RELATED WORK

There has been a rapid expansion in the usage of LEO satellite constellations for global connectivity. Previous research in LEO satellite connectivity has been dedicated to satellite network measurement [16, 27], edge computing in satellites [5, 24, 38], inter-satellite links [11, 13], satellite security [10, 25, 48], and satellite-ground station traffic scheduling [39, 45, 46]. However, most research presumes centralized control of satellite constellations and does not consider a decentralized approach.

The closest to our work is Mosaic [23], which focuses on sharing a single satellite by deploying networking stacks for multiple mobile network operators (MNOs) to enable direct-to-cell multi-tenancy. Such designs don't concern themselves with questions of coverage or constellation-level sharing, which are essential for multi-party LEO constellations.

The concept of sharing infrastructure has been explored in other networking domains, such as datacenters [15, 21] and distributed cellular networks [3, 12, 14, 26]. In datacenters, the practice of resource pooling allows multiple users to share computational and storage resources efficiently, optimizing utilization and reducing costs. There has been a similar approach in cellular networks [3]. It introduces a framework that enables MNOs to share base stations and backhaul networks, which reduces operational costs. While conceptually similar, as mentioned before, sharing satellite networks raise unique opportunities and challenges due to orbital dynamics.

6 CONCLUSION

We demonstrate that independent competing LEO constellations are prohibitively expensive, limit competition, and lead to unnecessary orbital occupancy. Instead, we provide a blueprint for shared multi-party LEO constellations. Participants in MP-LEO constellations cooperatively maximize coverage for themselves and others at much lower costs. We demonstrate that such designs can be robust to parties leaving the network or satellite failures. While our work highlights the promise of MP-LEO constellations, our designs are a work in progress and reveal a lot of open questions for the community to collectively tackle.

REFERENCES

[1] Amazon. 2022. AWS Ground Station: Easily control satellites and ingest data with fully managed Ground Station as a Service. https://aws.amazon.com/ground-station/.

[2] Amazon News. 2020. Amazon receives FCC approval for Project Kuiper satellite constellation. https://www.aboutamazon.com/news/company-news/amazon-receives-fcc-approval-for-project-kuiper-satellite-constellation. Accessed: June 26, 2024.

[3] Serhat Arslan, Ali Abedi, and Sachin Katti. 2023. d-Cellular: Trust-Free Connectivity in Decentralized Cellular Networks. In *2023 IEEE Future Networks World Forum (FNWF)*. 1–6. https://doi.org/10.1109/FNWF58287.2023.10520508

[4] Roger Baig, Ramon Roca, Felix Freitag, and Leandro Navarro. 2015. Guifi. net, a crowdsourced network infrastructure held in common. , 150–165 pages.

[5] Debopam Bhattacherjee, Simon Kassing, Melissa Licciardello, and Ankit Singla. 2020. In-orbit Computing: An Outlandish thought Experiment?. In *Proceedings of the 19th ACM Workshop on Hot Topics in Networks* (Virtual Event, USA) *(HotNets '20)*. Association for Computing Machinery, New York, NY, USA, 197–204. https://doi.org/10.1145/3422604.3425937

[6] Bloomberg. 2024. Pentagon Deal With Musk's Starlink in Ukraine Extended Six Months for $14 Million. https://www.bloomberg.com/news/articles/2024-06-13/pentagon-extends-deal-with-musk-s-starlink-in-ukraine-for-six-months. Accessed: June 26, 2024.

[7] CNN. 2022. SpaceX Starlink user terminals arrive in Ukraine, officials says. https://www.cnn.com/2022/02/27/business/starlink-activated-ukraine/index.html. Accessed: May 5, 2024.

[8] Bradley Denby and Brandon Lucia. 2020. Orbital edge computing: Nanosatellite constellations as a new class of computer system. In *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*. 939–954.

[9] Eutelsat Group. 2023. Satellite-as-a-Service: The Future of Satellite Network Communications. https://www.eutelsat.com/en/blog/satellite-as-a-service-the-future-of-satellite-network-communications.html.

[10] Giacomo Giuliari, Tommaso Ciussani, Adrian Perrig, and Ankit Singla. 2021. ICARUS: Attacking low Earth orbit satellite networks. In *2021 USENIX Annual Technical Conference (USENIX ATC 21)*. USENIX Association, 317–331. https://www.usenix.org/conference/atc21/presentation/giuliari

[11] Mark Handley. 2018. Delay is Not an Option: Low Latency Routing in Space. In *Proceedings of the 17th ACM Workshop on Hot Topics in Networks* (Redmond, WA, USA) *(HotNets '18)*. Association for Computing Machinery, New York, NY, USA, 85–91. https://doi.org/10.1145/3286062.3286075

[12] Shaddi Hasan, Amar Padmanabhan, Bruce Davie, Jennifer Rexford, Ulas Kozat, Hunter Gatewood, Shruti Sanadhya, Nick Yurchenko, Tariq Al-Khasib, Oriol Batalla, et al. 2023. Building flexible, low-cost wireless access networks with Magma. *GetMobile: Mobile Computing and*

Communications 27, 3 (2023), 40–47.

[13] Yannick Hauri, Debopam Bhattacherjee, Manuel Grossmann, and Ankit Singla. 2020. "Internet from Space" without Inter-satellite Links. In *Proceedings of the 19th ACM Workshop on Hot Topics in Networks* (Virtual Event, USA) *(HotNets '20)*. Association for Computing Machinery, New York, NY, USA, 205–211. https://doi.org/10.1145/3422604.3425938

[14] Helium Systems Inc. 2020. Helium A Decentralized Wireless Network. http://whitepaper.helium.com/.

[15] Benjamin Hindman, Andy Konwinski, Matei Zaharia, Ali Ghodsi, Anthony D Joseph, Randy Katz, Scott Shenker, and Ion Stoica. 2011. Mesos: A platform for {Fine-Grained} resource sharing in the data center. In *8th USENIX Symposium on Networked Systems Design and Implementation (NSDI 11)*.

[16] Liz Izhikevich, Manda Tran, Katherine Izhikevich, Gautam Akiwate, and Zakir Durumeric. 2024. Democratizing LEO Satellite Network Measurement. *Proc. ACM Meas. Anal. Comput. Syst.* 8, 1, Article 13 (feb 2024), 26 pages. https://doi.org/10.1145/3639039

[17] Min Jia, Ximu Zhang, Jintian Sun, Xuemai Gu, and Qing Guo. 2020. Intelligent Resource Management for Satellite and Terrestrial Spectrum Shared Networking toward B5G. *IEEE Wireless Communications* 27, 1 (2020), 54–61. https://doi.org/10.1109/MWC.001.1900238

[18] Kuiper Systems. 2020. Kuiper Systems, LLC Application for Authority to Deploy and Operate a Ka-band Non-Geostationary Satellite Orbit System. FCC, https://docs.fcc.gov/public/attachments/FCC-20-102A1_Rcd.pdf.

[19] John Liagouris, Vasiliki Kalavri, Muhammad Faisal, and Mayank Varia. 2023. SECRECY: Secure collaborative analytics in untrusted clouds. In *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)*. USENIX Association, Boston, MA, 1031–1056. https://www.usenix.org/conference/nsdi23/presentation/liagouris

[20] Ying-Chang Liang, Junjie Tan, Haonan Jia, Jintao Zhang, and Lian Zhao. 2021. Realizing Intelligent Spectrum Management for Integrated Satellite and Terrestrial Networks. *Journal of Communications and Information Networks* 6, 1 (2021), 32–43. https://doi.org/10.23919/JCIN.2021.9387703

[21] An-Dee Lin, Chung-Sheng Li, Wanjiun Liao, and Hubertus Franke. 2017. Capacity optimization for resource pooling in virtualized data centers with composable systems. *IEEE Transactions on Parallel and Distributed Systems* 29, 2 (2017), 324–337.

[22] Yehuda Lindell. 2020. Secure multiparty computation. *Commun. ACM* 64, 1 (2020), 86–96.

[23] Lixin Liu, Yuanjie Li, Hewu Li, Jiabo Yang, Wei Liu, Jingyi Lan, Yufeng Wang, Jiarui Li, Jianping Wu, Qian Wu, Jun Liu, and Zeqi Lai. 2024. Democratizing Direct-to-Cell Low Earth Orbit Satellite Networks. In *21st USENIX Symposium on Networked Systems Design and Implementation (NSDI'24)*. USENIX.

[24] Weisen Liu, Zeqi Lai, Qian Wu, Hewu Li, Qi Zhang, Zonglun Li, Yuanjie Li, and Jun Liu. 2024. In-Orbit Processing or Not? Sunlight-Aware Task Scheduling for Energy-Efficient Space Edge Computing Networks. In *IEEE International Conference on Computer Communications (INFOCOM'24)*. IEEE.

[25] Wei Liu, Yuanjie Li, Hewu Li, Yimei Chen, Yufeng Wang, Jingyi Lan, Jianping Wu, Qian Wu, Jun Liu, and Zeqi Lai. 2024. The Dark Side of Scale: Insecurity of Direct-to-Cell Satellite Mega-Constellations. In *The 45th IEEE Symposium on Security and Privacy (SP'24)*. IEEE.

[26] Zhihong Luo, Silvery Fu, Mark Theis, Shaddi Hasan, Sylvia Ratnasamy, and Scott Shenker. 2021. Democratizing cellular access with CellBricks. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*. 626–640.

[27] Sami Ma, Yi Ching Chou, Haoyuan Zhao, Long Chen, Xiaoqiang Ma, and Jiangchuan Liu. 2023. Network Characteristics of LEO Satellite Constellations: A Starlink-Based Measurement from End Users. In *IEEE INFOCOM 2023 - IEEE Conference on Computer Communications*. 1–10. https://doi.org/10.1109/INFOCOM53939.2023.10228912

[28] Yan Michalevsky and Yonatan Winetraub. 2017. SpaceTEE: Secure and Tamper-Proof Computing in Space using CubeSats. *CoRR* abs/1710.01430 (2017). arXiv:1710.01430 http://arxiv.org/abs/1710.01430

[29] Microsoft. 2023. CosmicBeats Simulator. https://github.com/microsoft/CosmicBeats-Simulator.

[30] Microsoft Azure. 2022. Azure Orbital Ground Station. https://azure.microsoft.com/en-us/products/orbital.

[31] National Aeronautics and Space Administration (NASA). [n. d.]. International Space Station. https://www.nasa.gov/international-space-station/.

[32] OneWeb. 2020. OneWeb Non-Geostationary Satellite System (LEO). FCC, https://fcc.report/IBFS/SAT-MPL-20200526-00062/2379706.pdf.

[33] PBS News. 2023. Elon Musk's refusal to provide Starlink support for Ukraine attack in Crimea raises questions for Pentagon. https://www.pbs.org/newshour/economy/elon-musks-refusal-to-provide-starlink-support-for-ukraine-attack-in-crimea-raises-questions-for-pentagon. Accessed: June 26, 2024.

[34] Reuters. 2021. Musk says may need $30 bln to keep Starlink in orbit. https://www.reuters.com/business/aerospace-defense/musk-sees-starlink-winning-500000-customers-next-12-months-2021-06-29/. Accessed: June 26, 2024.

[35] Reuters. 2022. Musk's Starlink connects remote Tonga villages still cut off after tsunami. https://www.reuters.com/world/asia-pacific/musks-starlink-connects-remote-tonga-villages-still-cut-off-after-tsunami-2022-02-23/.

[36] Reuters. 2024. Ukraine seeking action to stop Russian use of Starlink, minister says. https://www.reuters.com/world/europe/ukraine-seeking-action-stop-russian-use-starlink-minister-says-2024-02-19/. Accessed: June 26, 2024.

[37] SpaceX. 2016. SpaceX Non-Geostationary Satellite System. FCC, https://fcc.report/IBFS/SAT-LOA20161115-00118/1158350.pdf.

[38] Bill Tao, Om Chabra, Ishani Janveja, Indranil Gupta, and Deepak Vasisht. 2024. Known Knowns and Unknowns: Near-realtime Earth Observation Via Query Bifurcation in Serval. In *21st USENIX Symposium on Networked Systems Design and Implementation (NSDI 24)*. USENIX Association, Santa Clara, CA, 809–824. https://www.usenix.org/conference/nsdi24/presentation/tao

[39] Bill Tao, Maleeha Masood, Indranil Gupta, and Deepak Vasisht. 2023. Transmitting, Fast and Slow: Scheduling Satellite Traffic through Space and Time. In *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking* (Madrid, Spain) *(ACM MobiCom '23)*. Association for Computing Machinery, New York, NY, USA, Article 24, 15 pages. https://doi.org/10.1145/3570361.3592521

[40] TechCrunch. 2023. Loft Orbital is launching 'virtual missions' for developers wanting access to space. https://techcrunch.com/2023/11/29/loft-orbital-is-launching-virtual-missions-for-developers-wanting-access-to-space/.

[41] The Korea Bizwire. 2024. S. Korea to Develop LEO Satellite Communications System by 2030: Science Ministry. http://koreabizwire.com/s-korea-to-develop-leo-satellite-communications-system-by-2030-science-ministry/282001. Accessed: June 26, 2024.

[42] The New York Times. 2024. Why Taiwan Is Building a Satellite Network Without Elon Musk. https://www.nytimes.com/2024/03/14/business/taiwan-starlink-satellite.html?unlocked_article_code=1.gE0.yPkn.6yOsTiZIGhmg&smid=url-share. Accessed: May 5, 2024.

[43] The Spaceflight Now. 2024. SpaceX reaches nearly 6,000 Starlink satellites on orbit following Falcon 9 launch from Cape Canaveral. https://spaceflightnow.com/2024/05/12/live-coverage-spacex-to-reach-6000-starlink-satellites-on-orbit-following-falcon-9-launch-

from-cape-canaveral/. Accessed: June 26, 2024.

[44] The Verge. 2023. SpaceX's Starlink and other satellite internet providers are making light pollution worse for astronomers. https://www.theverge.com/2023/3/21/23649937/starlink-spacex-satellite-light-pollution-nature-astronomy. Accessed: June 26, 2024.

[45] Deepak Vasisht and Ranveer Chandra. 2020. A Distributed and Hybrid Ground Station Network for Low Earth Orbit Satellites *(HotNets '20)*. Association for Computing Machinery, New York, NY, USA, 190–196. https://doi.org/10.1145/3422604.3425926

[46] Deepak Vasisht, Jayanth Shenoy, and Ranveer Chandra. 2021. L2D2: low latency distributed downlink for LEO satellites *(SIGCOMM '21)*. Association for Computing Machinery, New York, NY, USA, 151–164. https://doi.org/10.1145/3452296.3472932

[47] James Richard Wertz, Wiley J Larson, Douglas Kirkpatrick, and Donna Klungle. 1999. *Space mission analysis and design*. Vol. 8. Springer.

[48] Johannes Willbold, Moritz Schloegel, Manuel Vögele, Maximilian Gerhardt, Thorsten Holz, and Ali Abbasi. 2023. Space Odyssey: An Experimental Software Security Analysis of Satellites. In *2023 IEEE Symposium on Security and Privacy (SP)*. 1–19. https://doi.org/10.1109/SP46215.2023.10351029

[49] YahooTech. 2024. SpaceX's Starlink internet service reaches milestone. https://www.yahoo.com/tech/spacex-reaches-milestone-starlink-internet-032052994.html. Accessed: June 26, 2024.