The Internet from Space, Reimagined: Leveraging Altitude for Efficient Global Coverage

Chris Misa
University of Oregon

Ramakrishnan Durairajan
University of Oregon and Link Oregon

Abstract

The Internet from Space has recently attracted renewed attention following technological developments that enable massive constellations of small satellites in low Earth orbit (LEO). While LEO satellite networks (LSNs) promise low-latency global connectivity, they face several fundamental challenges including the need for constant satellite replacement due to orbital decay, affecting environmental sustainability, and increasing congestion in orbital space from emerging players, heightening collision risks.

In this work, we propose to expand the LSN constellation design space by including use of altitude as a flexible design parameter to help solve the aforementioned challenges—i.e., by constructing constellations with orbits throughout the range implied by the classic LEO, medium Earth orbit (MEO), and geosynchronous orbit (GEO) designators. Although altitudes above LEO induce higher propagation latency, they also increase how much of the Earth's surface is visible to each satellite, thereby significantly reducing the total number of satellites required for global coverage. Building on this intuition, we provide an initial theoretical analysis of the tradeoffs enabled by orbital altitudes from LEO all the way up to GEO and conduct packet-level simulations demonstrating that MEO constellations can achieve present-day Internet latencies while using ∼19× fewer satellites and ∼14× fewer handovers than LEO.

CCS Concepts

• **Networks → Network design principles**; *Network performance evaluation*; *Mobile networks.*

ACM Reference Format:
Chris Misa and Ramakrishnan Durairajan. 2025. The Internet from Space, Reimagined: Leveraging Altitude for Efficient Global Coverage. In *3rd Workshop on LEO Networking and Communication (LEO-Net '25), September 8–11, 2025, Coimbra, Portugal.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3748749.3749092

1 Introduction

In recent years, several efforts in industry and academia alike have reignited global interest in the concept of providing Internet connectivity from space [20, 21, 33, 36, 38]. This resurgence is largely driven by the rapid development and deployment of low Earth orbit (LEO) satellite constellations (LSNs), which offer the promise of high-speed, low-latency Internet access to even the most remote and under-served regions of the world. Companies like SpaceX [16], OneWeb [9], SSST [35], Telesat [19], and Amazon [11] have made significant investments in building and launching such networks, with the goal of closing the digital divide and enabling universal connectivity [39, 41, 52, 56].

However, the low orbital altitudes (*e.g.,* typically between 500 and 2,000 km) that enable LSN's promise of low-latency, high-bandwidth Internet access [33] also present a key design trade-off. In particular, the coverage area of each satellite is significantly smaller than that of higher-altitude systems and atmospheric drag limits each satellite's orbital lifetime. As a result, global coverage and local bandwidth demands necessitate the deployment of "megaconstellations" of relatively cheap, disposable satellites. Given reduced per-satellite launch and manufacture costs as well as performance advantages of lower altitudes (*e.g.,* satisfying the <∼10 ms latency requirements for present-day direct-to-cell 5G service which, in turn, requires <∼1000 km altitude [44, 46]), such low-altitude megaconstellations make strong economic sense. However, a growing body of research raises questions about their scalability, sustainability, and long-term costs [5, 12, 30, 49, 53].

Although several recent efforts do tackle the issues of reducing the number of satellites required for high-performance satellite networking [24, 25, 40, 57], they miss the fundamental trade-offs enabled by orbital altitude. On the one hand, efforts that seek optimal points in the LSN design space [24, 40] are limited by the inherently large number of cheap, low-capacity satellites required to meet coverage and bandwidth demands (*e.g.,* ∼1k satellites to achieve continuous global coverage with minimal bandwidth). On the other hand, efforts that integrate higher orbital altitudes [25, 57] miss the nuanced interaction between altitude, coverage area, and propagation latency, focusing instead on routing and integration with the edge-cloud continuum.

To expand the design space, we envision a future where altitude is treated not as a fixed parameter, but as a critical and flexible component in constellation design [25, 26, 29]. Rather than narrow focus on LEO, medium Earth orbit (MEO), or GEO regimes, we envision satellite networks that fluidly blend a wide range of orbital altitudes while carefully balancing latency and bandwidth requirements with the effects of orbital mechanics, coverage radius, propagation latency, and environmental factors such as near-Earth radiation exposure [31, 42, 47]. As an initial step towards this vision, we present (to the best of our knowledge) the first analysis of the networking opportunities and challenges of deploying constellations at intermediate altitudes—*i.e.,* between traditional LEO, MEO, and GEO. We focus in particular on achieving global coverage because *(i)* the complex time-varying ground-tracks of LEO and MEO satellites make coverage of fixed geographic regions nearly impossible and *(ii)* given a constellation approach that satisfies global coverage, additional bandwidth demand can be addressed (up to a limit) by adding additional orbital planes with overlapping coverage

ACM ISBN 979-8-4007-2090-1/25/09
https://doi.org/10.1145/3748749.3749092

or (in the case of longer-lived, higher-altitude orbits) by increasing bandwidth capacity of each satellite (*e.g.,* by adding additional spot beams).

Our key insight is that despite their distinct challenges, altitudes just above traditional LEO open up the possibility to significantly reduce the total number of satellites required for global coverage while preserving reasonable network latency. Realizing this possibility demands inter-disciplinary theoretical and empirical exploration of the trade-offs involved. This work, in particular explores the networking implications of exploiting altitude as a powerful design variable rather than as a fixed constraint through the following three contributions. First, we present a theoretical model that quantifies the minimum number of satellites required to achieve global coverage as a function of altitude. Second, we approximate best-case networking performance metrics such as latency and coverage redundancy across a range of altitudes. Third, we empirically evaluate the impact of altitude on network behavior through realistic, simulation-based scenarios, providing insights that are grounded in operational parameters and practical constraints.

All scripts and NS3-based [8, 13] simulations are available at https://github.com/chris-misa/leveraging-altitude.

2 Background & Motivation

2.1 Goals of Satellite Networking

The key promise underlying the recent resurgence of interest in satellite networking is the ability of LEO satellite networks (LSNs) to potentially provide low-latency high-bandwidth connectivity to arbitrary regions of the Earth's surface. Previous generation satellite Internet efforts leveraged designs with small numbers of large satellites placed in high-altitude geosynchronous Earth orbits (GEO) that suffer from inherently high latency (*e.g.,* ∼100 ms one-way delay due to speed-of-light limit) and are incapable of supporting the low-latency requirements of modern web applications and real-time communication protocols. Present-day satellite Internet efforts embraced by the research community leverage designs with large numbers of small satellites placed in low-altitude Earth orbits (LEO) that satisfy lower latency requirements and raise a wide variety of novel technical challenges rooted in their dynamism.

As a *backhaul* for cellular-type networks, such connectivity enables expanding a wide range of latency-sensitive applications (*e.g.,* real-time communication, web browsing) to billions of "unconnected" users anywhere on Earth [20, 24]. Moreover, as a *backup* for terrestrial communication technologies, the global reach of LSNs can potentially improve resilience of critical Internet-based services—including emergency response coordination—in the face of multi-hazard risks (*e.g.,* earthquakes, wildfires, etc.) [52]. LSNs have also been proposed for specialized ultra-low latency applications, such as high-frequency trading, due to their potential for faster-than-fiber long-distance communication [33].

2.2 Limitations of Current Approaches

LSNs must be large. A fundamental challenge in LEO satellite networks (LSNs) is the limited coverage area of individual satellites. Due to their proximity to Earth, a single LEO satellite can only cover a relatively small footprint on the surface at any given time. This geographic coverage limitation necessitates deploying hundreds or thousands of satellites (i.e., mega-constellations) to ensure continuous global coverage.

Furthermore, the limited bandwidth of ground-to-satellite links (GSLs) introduces an additional constraint. In practice, a single satellite cannot satisfy the total demand of a given region, especially *e.g.,* in densely populated areas. This leads to the design requirement that multiple satellites must simultaneously cover each surface cell to meet bandwidth needs as well as provide redundancy. As a result, modern LSNs often incorporate overlapping satellite footprints, driving up the number of satellites needed.

Recent efforts have proposed non-uniform LEO constellations based on repeat ground-track orbits to optimize regional performance. For example, the approach described by [24] suggests selectively augmenting bandwidth in targeted areas through careful orbital planning. While such strategies may provide localized performance benefits, they do not eliminate the fundamental requirement for a larger number of satellites to ensure global coverage. As such, the actual cost savings and satellite reductions from these approaches remain uncertain in real-world conditions.

Satellites require continued replacements. In addition to the aforementioned physical and technical constraints, mega-constellations face significant sustainability challenges. Satellites in LEO are subject to atmospheric drag, which gradually degrades their orbits over time and ultimately causes reentry. Unlike higher-orbit satellites, LEO spacecraft have relatively short operational lifespans (e.g., typically on the order of 5 to 7 years) before they must be replaced [23, 45]. This impermanence introduces a steady-state requirement for LSNs: to maintain consistent network performance, operators must regularly launch new satellites to replace those that deorbit.

To quantify the impermanence, we analyze the replacement rate required to sustain a given satellite population using empirical lifetime estimates for real-world orbits. In particular, we use STELA [18] to propagate a sample of real-world StarLink orbits (altitude ∼550 km, inclination 53.0 degrees) obtained from Celestrak [1] to obtain a probability distribution (mean and 95-percent confidence interval based on Monte Carlo simulation) for the lifetime of these orbits. We then extrapolate these lifetime estimates to larger numbers of satellites and show in Figure 1a the annual number of replacement launches needed to maintain steady-state population (*e.g.,* in StarLink) as a function of the population size. Concurrently, Figure 1b estimates the amount of aluminum oxide (AlO) deposited into the mesosphere by reentering satellites, based on models from [30].

At the steady-state population of 1584 satellites, which corresponds to the target in a Federal Communications Commission (FCC) filing for one StarLink shell, the network would require the disposal of a median of 202 satellites per year. This process would release an estimated 6 metric tons of AlO annually. For context, this single network layer alone would contribute approximately 35% of the total AlO deposited from all satellite re-entries worldwide in 2022. Given that StarLink's full constellation includes multiple such layers with similar orbital parameters and "growing space race" across countries [4, 6, 7, 14], we believe the cumulative environmental impact will be substantial if not worse.

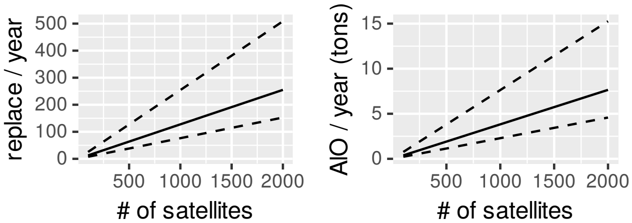

(a) No. of replacement events/year required to maintain a steady-state no. of satellites (x-axis). **(b) Corresponding amount of AlO deposited in the mesosphere by re-entries (based on [30]).**

Figure 1: Estimated environmental cost of maintaining a steady-state no. of satellites in one layer of StarLink.

Orbital space is getting overcrowded. As deployment of mega-constellations has accelerated (e.g., Starlink [16], Qianfan or "Thousand Sails" [35]), the resulting saturation of orbital slots introduces a number of critical operational risks [24, 34, 45, 49]. Primary among them is the increased probability of on-orbit collisions. To mitigate this risk, satellites must perform active collision avoidance maneuvers, which consume limited on-board fuel and thereby reduce their operational lifespan. These maneuvers also add complexity to inter- and intra-constellation management and require continuous coordination with other operators and regulatory bodies.

Beyond operational concerns, the growing presence of LEO satellites poses a serious threat to Earth-based astronomy. The reflective surfaces of satellites can interfere with optical observations, while their radio emissions can disrupt sensitive radio telescopes. Recent analyses highlight that unless mitigated, the scale of current and planned LEO deployments will significantly degrade the scientific value of both professional and amateur astronomical observations [12, 53]. This adds another dimension to the environmental and societal costs of scaling LSNs.

3 The Potential of Higher-Than-LEO Altitudes

The drawbacks of LEO mega-constellations discussed in the previous section can all be traced back to one critical defining design decision: the choice to focus *exclusively* on low-altitude Earth orbits. We investigate the decision of constellation altitude by first providing informal background on the near-earth environment above LEO, then considering how altitude impacts the networking potential of a single satellite and how these impacts shape the properties of a global-coverage constellation.

3.1 Near-Earth Space Radiation

Low-latency satellite networking's near exclusive focus on LEO is often informally justified by the need to avoid regions of near-Earth space known as the *Van Allen radiation belts* [31, 47]. Although these radiation belts are often described as fixed altitude ranges that must be avoided, their real-world structure and dynamics are far more complex. To illustrate, Figure 2 shows (electron) radiation intensity observed by the CIRBE spacecraft [42, 43] for two consecutive passes (~9 hours apart) through the Van Allen belts (on April 21, 2023).

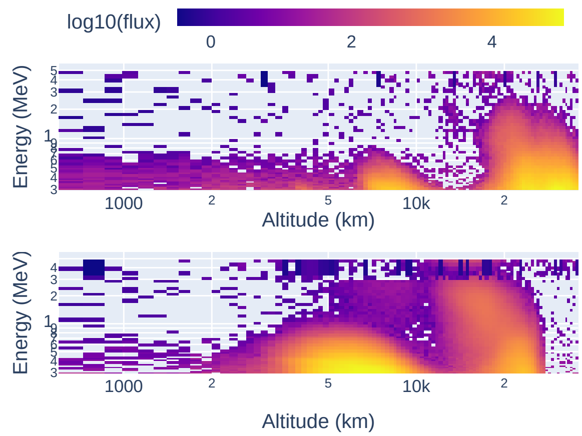

Figure 2: Example snapshots of radiation structure in near-Earth space between LEO and GEO altitudes.

Although two distinct regions of higher intensity radiation are apparent, they fluctuate significantly between observations, hinting at the wide-ranging dynamics apparent in CIRBE data [3]. Rather than clean "bands", near-Earth radiation is highly dynamic and at times (*e.g.,* during solar storms) crosses into regions typically considered outside of the bands (*e.g.,* the right-hand side of Figure 2 corresponds to GEO altitude; the South Atlantic Anomaly [15] persistently extends will into LEO altitudes below 500 km). For the design of satellite constellations for networking, this implies satellites will potentially need to deal with radiation *no matter their orbital altitude* and hence, the satellite networking community should not artificially limit the design space to narrow "safe" altitude bands.

3.2 Single-Satellite Impacts of Altitude

Given a more nuanced view of near-Earth radiation at different altitudes, we now turn to characterizing altitude's impact on networking, first, from the perspective of a single satellite, and later (§ 3.3) from the perspective of satellite constellation design.

Impact on network latency. The most widely-acknowledged impact of altitude on satellite networking is on latency induced by propagation delay. A lower-bound on this latency can be estimated by dividing the distance of a potential communication link by the speed of light. We visualize this relationship in Figure 3 by considering a single direction of a single ground-to-satellite radio link and showing the theoretical minimum speed-of-light latency (y-axis) for different altitudes (x-axis). We show this latency for both the shortest path, when the satellite is directly overhead the ground station (*i.e.,* at its zenith, solid blue), as well as a more realistic path, when the satellite is 30° above the ground station's horizon (dotted red). On the right of the figure, the relatively high altitudes required for GEO (~36,000 km) directly induce the unacceptably high network latency (>100 ms one-way) GEO has become associated with. On the left of the figure, the relatively low altitudes associated with LEO (*e.g.,* ~600 km) directly enable low-latency networking (*e.g.,* ~2 ms one-way).

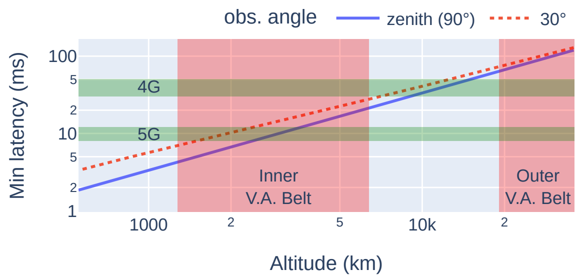

Figure 3: Impact of altitude on minimum one-way ground to satellite latency.

Between the extremes of LEO and GEO, Figure 3 illustrates a wide range of intermediate altitude and latency combinations. To contextualize, we draw horizontal bars (in green) indicating the approximate one-way latency ranges associated with 4G and 5G terrestrial radio links. We also draw vertical bars (in red) indicating the approximate average locations of the Van Allen radiation belts which increase operational costs by requiring addition shielding of sensitive radio components. Combining these ranges yields the insight that constellations at altitudes above the classic LEO range but below the inner Van Allen belt can still achieve lower one-way latency compared with 5G (*e.g.,* <10 ms) and a limited range of altitudes above the inner Van Allen belt can achieve lower one-way latency compared with 4G (*e.g.,* <40 ms).

Impact on Field-of-View. Beyond propagation latency, the next most important consideration in the design of satellite networks is the field-of-view (FoV), or area of the Earth's surface, that a single satellite can cover. Larger FoV implies a single satellite can cover more of the Earth's surface and hence fewer satellites are required for global coverage whereas smaller FoV implies the opposite.

Figure 4 shows FoV measured as the radius between the satellite's nadir and the circle defined by minimum observation angle (*i.e.,* elevation of the satellite above the ground station's horizon) for several minimum observation angles. Due to curvature of the Earth, lower-altitude satellites have smaller FoV whereas higher-altitude satellites have larger FoV. In particular (assuming a minimum observation angle of 30 degrees), the FoV of a LEO satellite at an altitude of ~600 km is ~800 km whereas the FoV of a GEO satellite is nearly 6000 km. (For comparison, an orthographic vantage point (*i.e.,* at infinite distance from the Earth) would cover half the Earth's circumference or ~20,000 km.)

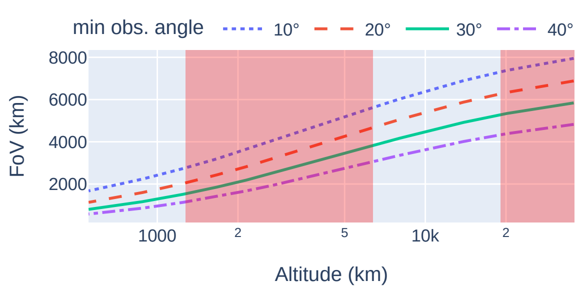

Figure 4: Impact of altitude on Field-of-View (FoV).

Figure 4 indicates that FoV increases significantly with constellation altitude. In particular, satellites above the inner Van Allen belt (again shown as vertical red regions) have over 2× higher FoV compared to typical LEO satellites below the belt. (Note that because we measure FoV as radius, this implies over 4× increase in the surface area covered.)

Impact on orbit lifetime. Finally, because LEO satellites operate in a near-Earth domain, their orbits gradually experience non-trivial decay due to atmospheric drag. The effects of atmospheric drag decrease with altitude so that higher-altitude orbits are inherently longer lived. In particular, orbits below ~1000 km may only last for $O(10)$ years where as higher orbits (above the influence of the atmosphere) last for $O(100-1000)$ years and are typically considered more-or-less permanent [23, 45].

3.3 Constellation-Wide Impacts of Altitude

We now integrate the above impacts of altitude on a single satellite's networking potential into the corresponding network-wide properties assuming a network that provides uniform global coverage (for latitudes up to its inclination).

Number of satellites required for global coverage. We first consider how designing higher altitude constellations can reduce the number of satellites required to provide global coverage. Assuming the common Walker-delta constellation geometry [27, 54] used by most recent LEO efforts and a fixed inclination of 65 degrees, we select orbital planes and satellites per plane in order to ensure complete coverage with non-zero overlap between satellites for handoffs. In particular, ascending nodes of orbital plans are separated by $2r\cos(\pi/4)$ and satellites in each plane are separated by $2r\cos(\pi/4)/\sin(i)$ where r is the FoV and i is the inclination.

Figure 5 shows the total number of satellites required by this approach (y-axis) as a function of altitude (x-axis) for several different minimum observation angles. We observe nearly exponential decrease in the number of satellites required as altitude increases (note the log axes). For example, a constellation at ~600 km requires ~1k satellites whereas a constellation at ~6000 km requires only ~60 satellites (assuming minimum observation angle of 30 degrees). This implies LEO altitudes require relatively large numbers of satellites to achieve minimal global coverage whereas constellations at higher altitudes achieve global coverage with far fewer satellites.

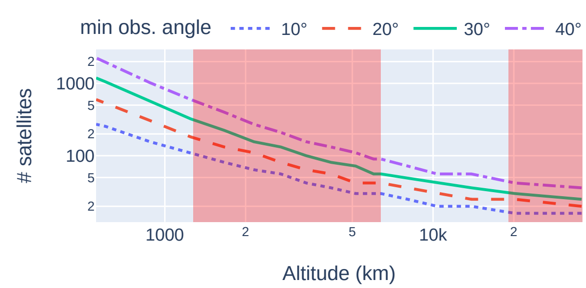

Figure 5: Minimum number of satellites required for global coverage as a function of constellation altitude.

Approximate end-to-end latency. Next, we consider the added latency overhead associated with higher altitude constellations. In particular, we assume per-hop overheads to be negligible and that the signal travels an arch along the orbital path directly between satellites. This approach approximates a low-bound on the latency achievable by a real-world constellation by replacing the jagged ISL paths required in topologies like +Grid [21, 36] or xGrid [48] with a smooth arch. After summing the total distance traveled in this way from source ground station, to the orbital altitude, and back to destination ground station, we again divide by the speed of light to estimate a theoretical lower-bound on the latency through a potential satellite network.

Figure 6 shows the approximated end-to-end latency for several different example terrestrial distances along the Earth surface. End-to-end latency remains below 100 ms for up to 10,000 km paths for constellations at altitudes of up to 6000 km indicating the potential for constellations above LEO to achieve usefully-low latency (*e.g.,* less than 100 ms). We also note that latency in such above-LEO constellations is significantly lower for shorter distances along the Earth surface indicating their potential use in applications like providing fiber-link backup (*e.g.,* over hundreds rather than thousands of km) in multi-hazard scenarios [52].

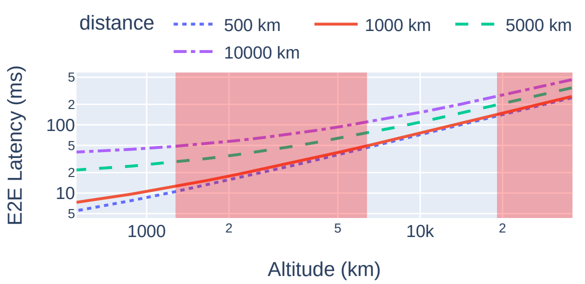

Figure 6: Approximate end-to-end latency for different terrestrial (great circle) distances.

4 Evaluation

To demonstrate practical utility of the theoretical and best-case estimates discussed above, we turn to packet-level simulation grounded in state-of-the-art models of orbital dynamics and link media.

4.1 Methodology

We simulate satellite networks using SNS3 [13, 50], a discrete event simulator (built on NS3 [8]) that models satellite orbital dynamics, inter-satellite links (ISLs), and (radio-based) ground-to-satellite links (GSLs). We construct example constellations at representative altitudes for GEO, MEO, and LEO using Walker-delta geometry as shown in Table 1. All constellations are designed to support global coverage up to 65° latitude with 30-degree minimum observation angle and have minimal overlap between coverage areas of adjacent satellites. For the LEO and MEO constellations we use static +Grid topologies with shortest-path routing. When making handovers, we only consider ascending satellites to avoid excessive path lengths due to the lack of (dynamic) links between satellites in opposite directions [51].

	Alt. (km)	# of Sats.	Inc.	# planes
GEO	35 786	5	∼ 0°	5
MEO	6000	56	65°	8
LEO	600	1054	65°	34

Table 1: LEO, MEO, and GEO constellation parameters.

For each constellation, we create two ground stations separated by 45° longitude at 41° latitude (∼3742 km great-circle distance apart). (Note that such connections present a worst-case workload for static +Grid topologies because of the need to leverage longitudinal ISLs between planes.) We model GSLs using slotted-ALOHA over Ka-band with custom antenna gain patterns based on constellation altitude. As in recent work [40], ISLs are assumed to be free-space optical with relatively high capacity compared to radio up/down links (*e.g.,* 100Mbps). Simulations and configuration files are available at https://github.com/chris-misa/leveraging-altitude.

4.2 Initial Results

We first compare the one-way latency experienced between the two ground stations in our simulation for the altitude domains represented by the constellations in Table 1. Figure 7a shows time-series of one-way latency for each constellation over one-hour of simulation time. To illustrate the root cause of latency fluctuation, we also show the number of ISL hops between ground stations in Figure 7b.

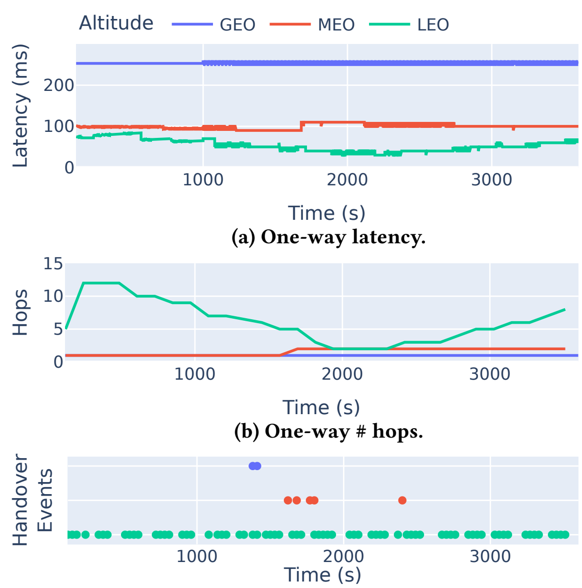

(a) One-way latency.

(b) One-way # hops.

(c) Handover events.

Figure 7: Comparison of one-hour simulation for example LEO, MEO, and GEO constellations.

In alignment with the results of § 3, we observe *(i)* GEO is hindered by high one-way latency (∼255 ms), *(ii)* LEO (at 600 km altitude) achieves low-latency (∼30 to ∼80 ms) but requires a large number of ISL hops (up to 12), and *(iii)* MEO (at 6000 km altitude) achieves slightly higher latency (∼100 ms) with only two or fewer ISL hops. On the one hand, for GEO this simulation result is very close to the theoretical minimum estimated in Figure 3. On the other hand, for LEO and MEO the increased latency compared to theoretical minimum is due to the non-direct ISL paths and radio-layer overheads of GSLs. However, our key observation is that the latency gap between LEO and MEO is relatively small, especially considering the ~19× fewer satellites used in MEO and relatively stable latency performance.

To illustrate the difference in dynamics between MEO and LEO constellations, in Figure 7c we visualize discrete handover events as points along the time axis. (Note that because SNS3 simulates GSLs with multiple "spot" beams for GEO satellites, we also observe occasional handover events between beams for the GEO constellation.) Over the hour of simulated time, we observed only five handovers in the MEO constellation compared to 70 handovers in the LEO constellation (~1 every minute). As a result of the increased FoV and decreased velocity of MEO (compared to LEO), this demonstrates MEO networks present significantly less challenging dynamics for networking compared to LEO.

5 Discussion and Outlook

Key Opportunities. The initial results presented in § 3 and § 4 establish the potential opportunities of considering a wider range of orbital altitudes in the design of networking satellite constellations. The first opportunity (enabled by increased FoV) is to reduce the total number of satellites required to provide similar near-global coverage compared to LEO. Although higher altitude MEO orbits do impose higher latency overheads, these overheads are significantly less compared to GEO and still fall within a reasonable range (*e.g.,* compared to present-day public Internet latencies over comparable distances [17, 22, 28, 55]). The second opportunity (enabled by longer orbit lifetimes) is to develop constellations of larger "heavy-weight" satellites with higher networking capacity (*e.g.,* more spot beams) and longer mission duration (*e.g.,* hundreds of years). This implies a constellation design and management model closer to GEO which has proven successful over time despite the challenges of high latency. The third opportunity (enabled by longer duration between handovers) is to simplify challenging networking requirements like routing and traffic engineering. We plan to explore joint constellation design (*e.g.,* LEO + MEO) by leveraging each of these opportunities in future work.

Key Challenges. Effectively leveraging the opportunities outlined above requires addressing several research and design challenges. The first challenge is in understanding the implications of radiation exposure as a function of orbit altitude and developing communication technologies that remain robust to radiation. This is not a new problem for satellite networks (*e.g.,* even at LEO altitudes, satellites receive considerable radiation exposure from the South Atlantic Anomaly [15, 32]) and real-world experimentation with higher-altitude orbits is already underway (*e.g.,* COSMOS-2553, designed specifically to fly at 2000 km directly in the lower Van Allen Belt [2, 10]). The second challenge is ensuring an optimum cost-benefit tradeoff in order to offset higher launch costs and larger MEO satellites (*e.g.,* with higher throughput enabled by more spot-beams). We again appeal to the GEO network service model here where careful long-term planning combined with relatively stable orbital dynamics provides a key leverage. Additional research efforts are required in this space to better understand the implications of critical constellation design parameters (*e.g.,* selection of Walker-delta constellation geometry vs. repeat ground-track geometry, combination of LEO and MEO altitudes [26, 29]) and their interaction with networking concerns (*e.g.,* exploration of dynamic topologies [21] enabled by the slower handover rate at higher altitudes). The third challenge is in optimizing ground-to-satellite radio protocols to enable fluid transition between the latency and power considerations of different altitudes. This can be seen as an extension or generalization of work already underway in the research community [44, 46] and the standards community [37].

Acknowledgments

We thank our anonymous reviewers for their constructive feedback. This work is supported by the Internet Society (ISOC) Foundation. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of ISOC.

References

[1] [n. d.]. CelesTrak. https://celestrak.org/. Accessed: 2024-12.
[2] [n. d.]. Cosmos-2553 - the first Neitron radar satellite. https://web.archive.org/web/20240522101825/https://russianforces.org/blog/2022/02/cosmos-2553_-_the_first_neitro.shtml. Accessed 2025-05-29.
[3] [n. d.]. Data Products - CIRBE. https://lasp.colorado.edu/cirbe/data-products/. Accessed 2025-05-29.
[4] [n. d.]. LEO Satellites in Orbit Skyrocket to Reach 42,600 Satellites by 2032 [...]. https://www.abiresearch.com/press/leo-satellites-in-orbit-skyrocket-to-reach-42600-satellites-by-2032-amid-growing-partnerships-and-competition?hsLang=en. Accessed 2025-05-29.
[5] [n. d.]. Mega-Constellations Pose Massive Risks to Shared Space Access. https://www.fierce-network.com/sponsored/mega-constellations-pose-massive-risks-to-shared-space-access. Accessed 2025-05-29.
[6] [n. d.]. Modern LEO Satellite Technologies and Space Race - Qorvo. https://www.qorvo.com/design-hub/blog/how-modern-leo-satellite-technologies-are-changing-the-space-race. Accessed 2025-05-29.
[7] [n. d.]. The New Space Race with China. https://www.palladiummag.com/2025/03/28/the-new-space-race-with-china. Accessed 2025-05-29.
[8] [n. d.]. ns-3 | a discrete-event network simulator for internet systems. https://www.nsnam.org/. Accessed 2025-05-29.
[9] [n. d.]. OneWeb. https://www.oneweb.world. Accessed 2025-05-29.
[10] [n. d.]. Opinion | What One Russian Satellite Tells Us About the Future of Nuclear Warefare - The New York Times. https://www.nytimes.com/interactive/2024/12/05/opinion/nuclear-weapons-space.html. Accessed 2025-05-29.
[11] [n. d.]. Project Kuiper. https://www.aboutamazon.com/what-we-do/devices-services/project-kuiper. Accessed 2025-05-29.
[12] [n. d.]. Satellite Licensing: FCC Should Reexamine Its Environmental Review Process for Large Constellations of Satellites. https://www.gao.gov/products/gao-23-105005. Accessed 2025-05-29.
[13] [n. d.]. SNS3. https://www.sns3.org/index.html. Accessed 2025-05-29.
[14] [n. d.]. The Space Rage is on. Connecting the Unconnected. | NWN. https://nwn.ai/blog/the-space-race-is-on-connecting-the-unconnected-with-low-earth-orbit-satellite-service/. Accessed 2025-05-29.
[15] [n. d.]. The spacecraft-killing anomaly over the South Atlantic. https://www.astronomy.com/space-exploration/the-spacecraft-killing-anomaly-over-the-south-atlantic/. Accessed 2025-05-29.
[16] [n. d.]. SpaceX Starlink. https://www.spacex.com/starlink. Accessed 2025-05-29.
[17] [n. d.]. Statistics for all Ark Monitors - CAIDA. https://www.caida.org/projects/ark/statistics/all_monitors/. Accessed 2025-05-29.
[18] [n. d.]. STELA | Connect by CNES. https://www.connectbycnes.fr/en/stela. Accessed: 2024-12.
[19] [n. d.]. Telesat Lightspeed LEO Network. https://www.telesat.com/leo-satellites/. Accessed 2025-05-29.
[20] [n. d.]. US agency approves T-Mobile, SpaceX license to extend covereage to dead zones. https://www.reuters.com/technology/us-agency-approves-t-mobile-spacex-license-extend-coverage-dead-zones-2024-11-26/. Accessed: 2025-05.
[21] Debopam Bhattacherjee and Ankit Singla. 2019. Network topology design at 27,000 km/hour. In *Proceedings of the 15th International Conference on Emerging Networking Experiments And Technologies*. 341–354.
[22] CJ Bovy, HT Mertodimedjo, Gerard Hooghiemstra, Henk Uijterwaal, and Piet Van Mieghem. 2002. Analysis of end-to-end delay measurements in Internet. In *Proc. of the Passive and Active Measurement Workshop-PAM*, Vol. 2002.
[23] EP Boykin and AR Mc Nair. 1966. *Earth orbital lifetime prediction model and program*. Technical Report.
[24] Yimei Chen, Lixin Liu, Yuanjie Li, Hewu Li, Qian Wu, Jun Liu, and Zeqi Lai. 2024. Unraveling Physical Space Limits for LEO Network Scalability. In *Proceedings of the 23rd ACM Workshop on Hot Topics in Networks*. 43–51.

[25] Yi Ching Chou, Xiaoqiang Ma, Feng Wang, Sami Ma, Sen Hung Wong, and Jiangchuan Liu. 2022. Towards sustainable multi-tier space networking for LEO satellite constellations. In *2022 IEEE/ACM 30th International Symposium on Quality of Service (IWQoS)*. IEEE, 1–11.

[26] Jean-Pierre Diris. [n. d.]. IRIS2: everything you need to know about this new european constellation. https://www.polytechnique-insights.com/en/columns/industry/iris2-everything-you-need-to-know-about-this-new-european-constellation/. Accessed 2025-05-29.

[27] John E Draim. 1987. A common-period four-satellite continuous global coverage constellation. *Journal of Guidance, Control, and Dynamics* 10, 5 (1987), 492–499.

[28] Ramakrishnan Durairajan, Sathiya Kumaran Mani, Paul Barford, Robert Nowak, and Joel Sommers. 2018. Timeweaver: Opportunistic one way delay measurement via ntp. In *2018 30th International Teletraffic Congress (ITC 30)*, Vol. 1. IEEE, 185–193.

[29] Directorate-General for Defence Industry European Commission and Space. [n. d.]. IRIS2 | Secure Connectivity - European Commission. https://defence-industry-space.ec.europa.eu/eu-space/iris2-secure-connectivity_en. Accessed 2025-05-29.

[30] José P Ferreira, Ziyu Huang, Ken-ichi Nomura, and Joseph Wang. 2024. Potential ozone depletion from satellite demise during atmospheric reentry in the era of mega-constellations. *Geophysical Research Letters* 51, 11 (2024), e2024GL109280.

[31] N Yu Ganushkina, I Dandouras, YY Shprits, and J Cao. 2011. Locations of boundaries of outer and inner radiation belts as observed by Cluster and Double Star. *Journal of Geophysical Research: Space Physics* 116, A9 (2011).

[32] Kirolosse M Girgis, Tohru Hada, Shuichi Matsukiyo, and Akimasa Yoshikawa. 2022. Radiation analysis of LEO mission in the South Atlantic anomaly during geomagnetic storm. *IEEE Journal of Radio Frequency Identification* 6 (2022), 292–298.

[33] 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*. 85–91.

[34] Daniel Jang, Andrea D'Ambrosio, Miles Lifson, Celina Pasiecznik, and Richard Linares. 2022. Stability of the LEO Environment as a Dynamical System. In *Advanced Maui Optical and Space Surveillance Technologies (AMOS) Conference*.

[35] Andrew Jones. [n. d.]. China's Qianfan Megaconstellation Aims at SpaceX's Starlink. https://spectrum.ieee.org/satellite-internet. Accessed 2025-05-29.

[36] Simon Kassing, Debopam Bhattacherjee, André Baptista Águas, Jens Eirik Saethre, and Ankit Singla. 2020. Exploring the" Internet from space" with Hypatia. In *Proceedings of the ACM Internet Measurement conference*. 214–229.

[37] Joern Krause. [n. d.]. Non-Terrestrial Networks (NTN). https://www.3gpp.org/technologies/ntn-overview. Accessed 2025-07.

[38] Zeqi Lai, Hewu Li, Yangtao Deng, Qian Wu, Jun Liu, Yuanjie Li, Jihao Li, Lixin Liu, Weisen Liu, and Jianping Wu. 2023. StarryNet: Empowering researchers to evaluate futuristic integrated space and terrestrial networks. In *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)*. 1309–1324.

[39] Zeqi Lai, Weisen Liu, Qian Wu, Hewu Li, Jingxi Xu, Yibo Wang, Yuanjie Li, and Jun Liu. 2024. SpaceRTC: Unleashing the Low-latency Potential of Mega-constellations for Wide-Area Real-Time Communications. *IEEE Transactions on Mobile Computing* (2024).

[40] Zeqi Lai, Yibo Wang, Hewu Li, Qian Wu, Qi Zhang, Yunan Hou, Jun Liu, and Yuanjie Li. 2024. Your Mega-Constellations Can Be Slim: A Cost-Effective Approach for Constructing Survivable and Performant LEO Satellite Networks. In *IEEE INFOCOM 2024-IEEE Conference on Computer Communications*. IEEE, 521–530.

[41] Jihao Li, Hewu Li, Zeqi Lai, Qian Wu, Yijie Liu, Qi Zhang, Yuanjie Li, and Jun Liu. 2024. Satguard: Concealing endless and bursty packet losses in leo satellite networks for delay-sensitive web applications. In *Proceedings of the ACM Web Conference 2024*. 3053–3063.

[42] Xinlin Li. 2024. Unveiling energetic particle dynamics in the near-Earth environment from CubeSat missions. *AGU Advances* 5, 3 (2024), e2024AV001256.

[43] Xinlin Li, Zheng Xiang, Yang Mei, Declan O'Brien, David Brennan, Hong Zhao, Daniel N Baker, and Michael A Temerin. 2025. A new electron and proton radiation belt identified by CIRBE/REPTile-2 measurements after the magnetic super storm of 10 May 2024. *Journal of Geophysical Research: Space Physics* 130, 2 (2025), e2024JA033504.

[44] Yuanjie Li, Hewu Li, Wei Liu, Lixin Liu, Yimei Chen, Jianping Wu, Qian Wu, Jun Liu, and Zeqi Lai. 2022. A case for stateless mobile core network functions in space. In *Proceedings of the ACM SIGCOMM 2022 Conference*. 298–313.

[45] J-C Liou and Nicholas L Johnson. 2008. Instability of the present LEO satellite populations. *Advances in Space Research* 41, 7 (2008), 1046–1053.

[46] Lixin Liu, Yuanjie Li, Hewu Li, Jiabo Yang, Wei Liu, Jingyi Lan, Yufeng Wang, Jiarui Li, Jianping Wu, Qian Wu, et al. 2024. Democratizing direct-to-cell low earth orbit satellite networks. *GetMobile: Mobile Computing and Communications* 28, 2 (2024), 5–10.

[47] BH Mauk, NJ Fox, SG Kanekal, RL Kessel, DG Sibeck, and A Ukhorskiy. 2013. Science Objectives and Rationale for the Radiation Belt Storm Probes Mission. *Space Science Reviews* 179, 1-4 (2013), 3–27.

[48] Joseph Mclaughlin, Jee Choi, and Ramakrishnan Durairajan. 2023. × grid: A location-oriented topology design for LEO satellites. In *Proceedings of the 1st ACM Workshop on LEO Networking and Communication*. 37–42.

[49] Amy Mehlman, John P. Janka, and Mark A. Sturza. [n. d.]. Averting Environmental Risks In The New Space Age. https://fas.org/publication/averting-environmental-risks-in-the-new-space-age/. Accessed 2025-05-29.

[50] Jani Puttonen, Budiarto Herman, Sami Rantanen, Frans Laakso, and Janne Kurjenniemi. 2015. Satellite network simulator 3. In *Workshop on Simulation for European Space Programmes (SESP)*, Vol. 24. 26.

[51] Xiaoxin Qi, Bing Zhang, Zhiliang Qiu, and Ling Zheng. 2021. Using inter-mesh links to reduce end-to-end delay in walker delta constellations.. In *IEEE Communications Letters)*, Vol. 25. 2070–2074. Issue 9.

[52] Aleksandr Stevens, Blaise Iradukunda, Brad Bailey, and Ramakrishnan Durairajan. 2024. Can LEO Satellites Enhance the Resilience of Internet to Multi-hazard Risks?. In *International Conference on Passive and Active Network Measurement*. Springer, 170–195.

[53] Constance Walker, Jeffrey Hall, Lori Allen, Richard Green, Patrick Seitzer, Tony Tyson, Amanda Bauer, Kelsie Krafton, James Lowenthal, Joel Parriott, et al. 2020. Impact of satellite constellations on optical astronomy and recommendations toward mitigations. *Bulletin of the American Astronomical Society* 52, 2 (2020).

[54] John G Walker. 1984. Satellite constellations. *Journal of the British Interplanetary Society* 37 (1984), 559.

[55] Bo Zhang, TS Eugene Ng, Animesh Nandi, Rudolf Riedi, Peter Druschel, and Guohui Wang. 2006. Measurement based analysis, modeling, and synthesis of the internet delay space. In *Proceedings of the 6th ACM SIGCOMM conference on Internet measurement*. 85–98.

[56] Jinwei Zhao and Jianping Pan. 2024. Low-Latency Live Video Streaming over a Low-Earth-Orbit Satellite Network with DASH. In *Proceedings of the 15th ACM Multimedia Systems Conference*. 109–120.

[57] Xiangming Zhu and Chunxiao Jiang. 2022. Delay optimization for cooperative multi-tier computing in integrated satellite-terrestrial networks. *IEEE Journal on Selected Areas in Communications* 41, 2 (2022), 366–380.