How to Dream in Space (Without Getting Burned): A Perspective on LEO Networking Research

Chris Misa
University of Oregon

Zander Chown
University of Oregon

Ramakrishnan Durairajan
University of Oregon & Link Oregon

Abstract

Low-Earth orbit (LEO) satellite networks (LSNs) have become one of networking's most exciting new frontiers. Yet much of the field rests on a fragile foundation: *simulations and modeling assumptions that are rarely scrutinized against the physical, operational, and environmental realities of the near-Earth space environment.* As researchers increasingly move beyond today's commercial deployments to propose entirely new satellite constellations custom-built for specific applications, we risk drifting into physical fiction, producing designs that satisfy communication objectives while violating fundamental constraints of orbital mechanics, spacecraft operations, or the space environment. We argue that this is not an isolated problem, but an emerging methodological gap in LSN research. In this perspective paper, we illustrate through representative examples from recent literature, argue for a stronger interface between the networking and space science communities, and propose an initial set of principles to ensure future LSN research is both ambitious and physically realizable.

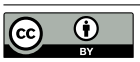
CCS Concepts

• **Networks** → **Network design principles**; *Network simulations.*

Keywords

Space-based computing infrastructure; LEO satellite networks; space mission analysis and design

ACM Reference Format:
Chris Misa, Zander Chown, and Ramakrishnan Durairajan. 2026. How to Dream in Space (Without Getting Burned): A Perspective on LEO Networking Research. In *The 25th ACM Workshop on Hot Topics in Networks (HotNets '26), November 16–17, 2026, Salt Lake City, UT, USA.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3847290.3859968



This work is licensed under a Creative Commons Attribution 4.0 International License.
HotNets '26, Salt Lake City, UT, USA
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-3029-0/2026/11
https://doi.org/10.1145/3847290.3859968

1 Introduction

Low-Earth orbit (LEO) satellite networks (LSNs) have emerged as a technology poised to enable novel solutions to a variety of persistent Internet connectivity challenges. The potential of LSNs to provide low-latency, high-bandwidth, and reliable connectivity *nearly anywhere in the word* implies a tantalizing set of high-impact applications including direct-to-consumer wireless in regions where terrestrial infrastructure is economically infeasible to realize [4, 44, 55], high-performance uplink for long-distance transportation (*e.g.,* airlines, shipping, trucking) [8, 14, 40], and reliable backup in regions affected by war or natural disasters [5, 52, 60]. Accordingly, networking research efforts have tackled questions of how best to leverage inter-satellite topologies and routing [7, 13, 25, 35, 56, 60], how to optimize radio networking technologies for satellite connectivity [38, 41–43, 57], and how to design constellation geometries that put satellites where and when they are most needed [20, 32, 37, 50].

Despite the promising applications, LSNs force the networking research community to reckon with a fundamental tension between creatively dreaming up space-based solutions to networking problems and pragmatically ensuring these solutions remain grounded in real-world constraints. Unlike terrestrial infrastructure, space-based platforms are constantly in motion, subject to complex physical constraints, and exposed to a harsh radiation environment. Networking research that fails to demonstrate a feasible path to realization through this maze of space-specific constraints should be dismissed as unrealistic.

It is typically also impractical for networking researchers to directly run experiments atop the space-based platforms leveraged by LSNs, leading to reliance on intuition and simulation-based evaluation techniques. Moreover, limited publicly available data about the very small number of space-based systems actually deployed in the real world (typically large-capital, for-profit operations like Starlink [3]) does not provide a sufficiently rich sample of the full design space for true scientific study of the complex tradeoffs involved. We argue that this has created *a methodological blind spot in LSN research:* proposals are often evaluated primarily through the lens of networking performance, with comparatively little attention paid to whether the underlying space-based infrastructure is physically valid.

Ironically, the networking community has largely navigated this challenge by tightly coupling research proposals to carefully gleaned facts about technologies deployed in present-day systems such as Starlink.[1] This strategy builds LSN research "directly atop" existing constellations by leveraging publicly available artifacts, such as ephemeris data (*e.g.,* [1]) and FCC filings (*e.g.,* [21]), that capture fragmentary facts about deployed systems. Rather than explicitly reasoning about the physical constraints of the near-Earth space environment or design tradeoffs in the engineering of satellites themselves (*e.g.,* the number of ISLs a satellite can form as discussed in § 2.1), proposals inherit their realism from the systems they emulate. Although it has proven remarkably effective for studying networking questions within the context of today's deployments, this strategy begins to break down rapidly once research moves beyond existing constellations and starts projecting networking concerns into design of the underlying space infrastructure itself.

Recently, the breakdown of this strategy has become strikingly evident through a new thread of research that investigates how to optimize constellation designs for networking requirements [20, 32, 37, 50]. In particular, motivated by recent results (from outside the networking community) that demonstrate dire environmental consequences of unchecked proliferation of massive-scale satellite constellations [2, 22, 48], these efforts primarily ask how to minimize the total number of satellites required to still meet the targeted networking goals. Addressing this question, by definition, requires proposing constellations that do not yet exist. This, in turn, requires networking proposals to account for a wide range of unexpected tradeoffs, constraints, degrees of freedom, and other factors inherent in the design and implementation of space-based infrastructure.

Despite initial promising results, many proposals in this new thread of networking-informed constellation design (published in top-tier networking venues) violate basic facts about space-based infrastructure, providing *illusory solutions* to complex interdisciplinary problems. As an illustrative example, a recent SIGCOMM publication [37] demonstrates how to reduce the number of satellites required for an LSN like Starlink by up to $\sim 7.9\times$, but, in doing so, requires a physically infeasible constellation where satellite trajectories *run into each other* on a regular basis (*e.g.,* every day) with high probability (as we demonstrate in § 2.3).

Motivated by these limitations, the primary objectives of this perspective paper are to explicitly bring the fundamental tension—*i.e.,* between real-world applicability and creative freedom—to the attention of the networking research community and to outline a set of principles that can serve as an initial interface to bridge the tension in future research. We seek to establish a "knowledge hub"[2] and set of methodological checks (*e.g.,* simulation platforms, review criteria) that can free future LSN research from the constraints of present-day deployments while simultaneously ensuring networking proposals are consistent with physical (non-networking) constraints. As a first step, we propose in § 3 a non-exhaustive set of critical principles that future LSN research should adhere to, especially efforts in directions that inherently extend beyond the set of possibilities captured by present-day deployments like Starlink. These initial principles imply that LSN research cannot continue solely in the context of the networking community (based on carefully gleaned facts from large corporations like SpaceX), but requires interdisciplinary collaboration across networking research, the space physics community, and present-day LSN operators like Starlink.

What makes this work both timely and provocative is our observation that LSN research has reached a critical inflection point. As the community moves beyond studying existing constellations toward inventing entirely new space-based infrastructures (*e.g.,* demand-tailored networking constellations [37, 50], space-based CDNs and VNFs [23, 70], etc.), realism can no longer be inherited from deployed systems or established through networking expertise alone. Instead, advancing the field (or determining if new proposals are at all feasible in the first place [10, 53]) requires a new research paradigm built on explicit connections between networking, orbital mechanics, spacecraft engineering, and space physics. By establishing an initial collection of such connections, we aim to enable a future where networking researchers can explore ambitious constellation designs, without being constrained by today's deployments, while avoiding getting "burned" by the trap of proposals that are elegant in simulation but impossible in space.

Disclaimers. Our focus in this work is on the interface between networking system design and the realities of space-based infrastructure. We explicitly do not consider the wide range of higher-level considerations (*e.g.,* authentication and trust, virtualization and resource sharing, etc.) that accompany the emergence of space-based infrastructure as a platform for Internet connectivity and other services beyond forwarding. These other considerations may themselves require parallel effort to organize defining principles, but such effort is beyond the scope of this present work.

We also emphasize that, although we cite specific prior works to illustrate our arguments, we do not intend to single

[1] In fact, this has led to a *whole new* line of research dedicated to inferring the system constraints, design decisions, and operational policies of real-world LSN deployments [11, 18, 51, 61, 69].

[2] To enable growth beyond this paper, we also establish an open-source online resource at https://spacenet-hub.github.io/.

these works out or to detract from the important contributions they have made to the field. The issues we identify here are (to the best of our knowledge) prevalent across the field (including in our own previously published works) and our focus is on enabling future research and real-world impact.

2 Example Challenges of LSN Research

To motivate the need for a new "knowledge hub" for LSN research, we consider how gleaning from limited publicly-available data about present-day LSN deployments fundamentally limits LSN research (§ 2.1), the challenges of building intuition about the complex dynamics of satellite-based infrastructure required to go beyond these limits (§ 2.2), and then an example of the unfortunate results when the community ignores these intuitions (§ 2.3).

2.1 Carefully-Gleaned Facts: How Many "Ports" Should a Satellite Have?

A key aspect of LSN's promise to provide connectivity solutions in scenarios where terrestrial technologies struggle is the ability of satellites to form direct peer-to-peer links with other satellites. By leveraging high-capacity free-space optics, inter-satellite links (ISLs) can achieve bandwidths comparable to terrestrial fiber (*e.g.,* up to 200 Gbps per ISL [58]) thus enabling high-capacity, low-latency connectivity between nearly any two points on the Earth's surface. Significant research effort has explored the potential benefits and challenges of ISL-based networking in LSNs [7, 11–13, 25, 35, 37, 66]. In this research direction, satellites essentially become space-based routers with (wireless) "ports" for access links and (optical) "ports" for forming ISLs.

As with terrestrial routers, a key engineering consideration is the number of links a satellite can form with other satellites, or the number of (non-access) ports. In particular, each ISL requires a distinct laser or radio transceiver on both connected satellites. Intuitively, adding more transceivers on the one hand enables more sophisticated topologies (*e.g.,* with higher redundancy or more direct paths), but on the other hand increases a satellite's mass, cost, and complexity (and hence the capital and operational expenditure of the LSN). Accordingly, constellation design decisions should exploit a tradeoff between adding more ISL transceivers per satellite to increase connectivity potential while minimizing the associated capital, operational, and environmental costs.

Given the inherent tradeoff in how many ISL-forming "ports" each satellite is equipped with, a natural direction for LSN research would be to first characterize this tradeoff, then leverage it in the design of constellation geometry, topology, and routing in order to meet specific networking objectives (*e.g.,* latency, throughput, path redundancy) with minimal overheads.

Unfortunately, the LSN research community has taken a rather different direction, following breadcrumbs inferred from limited and incomplete information about planned and existing LSN deployments on an implicit search for the "correct" number of ISLs per satellite. As shown in Table 1, the community developed a rough consensus that four ISL transceivers per satellite is an appropriate choice. Some proposals assume more transceivers are available (up to 12 per satellite [12]). Most significantly, this consensus lead to the development of the +Grid [7] ISL topology which was then assumed as a default for constellations like Starlink in subsequent research. Proposals for going beyond +Grid also typically assume four transceivers per satellite as a minimum specification [7, 47].

Year	Proposal	ISLs / Sat.	Breadcrumbs
1998	Werner et al. [66]	4	
2002	MLSR [6]	4	
2018	Handley et al. [25]	5	← SpaceX FCC filing mentions 4 potential ISL transceivers per satellite [26].
2019	Motif [7]	4	
2020	Hypatia [28]	4	
	Dai et al. [19]	4	
2021	Leyva et al. [35]	4	
	Chen et al. [13]	4	
	Chaudhry et al. [11]	4	
	Qi et al. [54]	5	
2022	SMTSN [17]	5	
	OWSN [12]	12	
	Gao et al. [23]	4	
2023	Liu et al. [43]	4	
	pLEO [56]	4	
	StarryNet [30]	4	
	xGrid [47]	4	
2024	SatGuard [36]	4	← Starlink progress report [58] explicitly mentions three ISL transceivers per satellite.
	SpaceRTC [31]	4	
	MegaReduce [33]	4	
	MTDA [68]	4	
2025	SaTE [67]	6	
	TinyLEO [37]	3	

Table 1: A sample of LSN publications showing how many ISL transceivers were assumed to be realistic.

Although supporting assumptions by referencing existing systems is generally a solid research practice, in this case it has lead to a somewhat futile quest to figure out existing deployments rather than investigating the networking implications of the underlying tradeoff space. For example, as pointed out by [26], the rough consensus holding until around 2025 was based on speculation around a first-generation Starlink proposal that was *never actually deployed*. Present-generation Starlink satellites (V2 Mini) actually have *three* laser-based ISL transceivers [58, 59]. This implies that the vast body of LSN research assuming +Grid or other 4-transceiver topologies atop the Starlink constellation would actually require replacing *all* existing

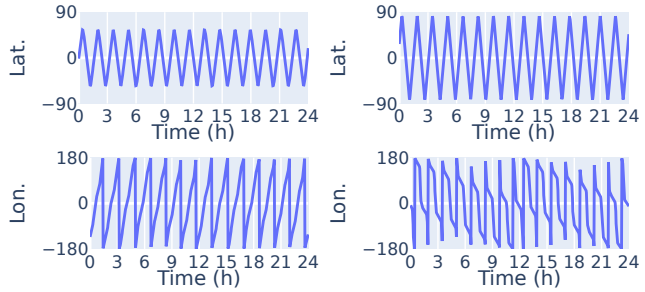
Starlink satellites. Moreover, Starlink did not actually start launching satellites equipped with laser ISL tranceivers until 2021 [27, 71] and it is unclear exactly how many ISLs are launched and how they are currently being used.[3] Recent research (*e.g.,* TinyLEO [37]) appears to be "correcting" this trend towards what Starlink has now published (3 per satellite), but still misses the underlying design tradeoff.

Ultimately, the question of how many links should each satellite be able to form requires interdisciplinary investigation of the tradeoff space induced by the choice of how many ISL transceivers are deployed in which particular LSN orbits. Just like how designers of terrestrial routers engineer products based on a combination of networking demands and physical constraints, so to must designers of space-based systems engineer satellites based on a combination of networking demands and physical constraints. Rather than limiting ourselves to unreliable estimates of existing deployments, the networking community should investigate the number of ISLs each satellite forms as a critical tradeoff. For example, we can adopt and refine an abstract cost model in collaboration with space technology engineers and work to identify potential "sweet spots" such that the number of ISLs is most effective for satisfying networking objectives while minimizing constellation-wide costs.

2.2 Where and When Are Your Routers?

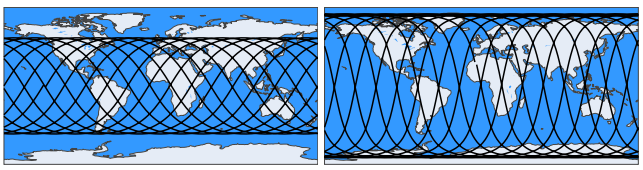
Beyond uncertainty about realistic capabilities of networking satellites, the next most cited (and yet still poorly understood) challenge in LSN research is the "extreme mobility" of satellites compared to traditional networking infrastructure. Understanding the extreme mobility is particularly critical for recent efforts that go beyond the designs and ephemeris[4] of real-world deployments (*e.g.,* [32, 37, 50]). Unlike terrestrial networks, where nodes are largely stationary and links are relatively stable, satellites move at high velocities along continuously changing orbital trajectories. As a result, network topology evolves rapidly over time, leading to intermittent connectivity, frequent handovers, and highly dynamic routing conditions.

Developing a comprehensive intuition for satellite mobility is particularly challenging because it requires not only understanding orbital dynamics, but also how the trajectories of satellites relate to the trajectories of points on the (spinning) Earth surface. To illustrate this challenge, we consider three different ways to visualize two orbits (a direct orbit at moderate inclination of 55° and a retrograde orbit at high inclination of 97.7°) over the course of one day. (For simplicity, we tune the altitude to the spin of the Earth such that both orbits are repeat ground-track (RGT) [63].) First, in Figure 1, we show the latitude and longitude trace on the Earth surface of each orbit over time. Second, Figure 2 shows the "ground track" of the same orbits—*i.e.,* the set of points on the Earth surface, visualized geographically, that the orbits pass above as time progresses. Finally, Figure 3 shows the positions traced by the same orbits in absolute three-dimensional space around the Earth.



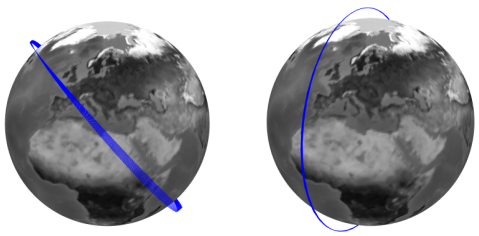
(a) Inclination 55° **(b) Inclination 97.7°**

Figure 1: Example time series of satellite positions in ~550 km altitude RGT orbits.



(a) Inclination 55° **(b) Inclination 97.7°**

Figure 2: "All-times" visualization of the same trajectories shown in Figure 1 on a equirectangular projection.



(a) Inclination 55° **(b) Inclination 97.7°**

Figure 3: "Fixed-globe" visualization of the same trajectories shown in Figure 1 in 3-dimensional space.

Each visualization captures different intuitions about satellite mobility, illustrating the challenges of "dreaming" accurately about the positions of LSN infrastructure, and hence

[3]The most likely possibility is that ISLs in Starlink serve to enable an "extended bent-pipe" configuration rather than as a global backbone network [37, 41, 51].

[4]Ephemeris are publicly-available records of where the satellites in existing deployed constellations are and their current trajectories.

the ability of this infrastructure to meet Earth-based networking demand. Figure 1 captures most of the relevant spatiotemporal dynamics, however, it is nearly impossible to connect the satellite's position to specific geographic concepts and adding the trajectories of more than a single satellite quickly yields an incomprehensible plot. Figure 2 explicitly connects the satellite's trajectory over time to particular geographies, however, it is nearly impossible to reason about *when* the satellite is over (and hence in-range) of these geographies and again adding more than a single satellite quickly makes the plot incomprehensible. Figure 3 is in many ways the simplest representation (*e.g.,* circular orbits appear more-or-less as circles) and allows visualizing multiple independent satellite orbits simultaneously (*e.g.,* see Figure 11 in [28]), however, it provides a false sense of connection to geography because in reality the satellite is only at one point on its orbit at a time and the Earth spins continuously underneath it. Moreover, space-time relationships gradually change as the Earth orbits the Sun making it challenging to reason about the local time-of-day at the points the satellite passes over.

Ultimately, the challenge of extreme mobility requires LSN researchers—and, critically, reviewers of LSN proposals—to develop accurate mental models of orbital dynamics by carefully inspecting multiple forms of visualization (and possibly animations and/or other interpretive tools). The form of mental gymnastics required is similar to the classic computer science exercise of thinking through all the possible interleavings of parallel threads of execution. In particular, imagining the trajectories of satellites in an LSN requires considering the concurrent motions of the Earth around its axis, the satellites around their orbits, the orbits themselves (as they are dragged around by Earth's oblateness among other perturbations), and the motion of the whole system around the Sun (which critically determines what time it is at points on the Earth surface). The field of orbital dynamics has already developed some relevant proofs and properties (*e.g.,* repeat ground-track orbits, Sun-synchronous orbits, etc.) that LSN researchers can use to simplify this and some initial proposals foreshadow the potential of this direction [63].

2.3 When Your Routers Are in the Same Place at the Same Time

To illustrate the extent of what can go wrong when LSN proposals miss a constraint imposed by satellite's extreme mobility, we consider a constellation geometry proposed by TinyLEO [37] that reduces the number of satellites required by leveraging tight matching of orbital parameters with terrestrial networking demand. TinyLEO's constellation geometry inherently departs from the common Walker-delta model [64, 65] and hence cannot be evaluated in terms of well-known constellation constraints or existing real-world ephemeris.

To demonstrate how TinyLEO actually violates basic physical constraints, we estimate the Euclidian distance (in three-dimensional space) between all satellites every ~5 minutes over the course of one day in the proposed constellation. Figure 4 shows where two satellites pass within 10 km of each other projected down onto latitude, longitude, time coordinates. (Note that because of numeric error in orbit propagation algorithms, 10 km is considered an informal minimum safe distance between objects in space operations [15, 16, 39].) As shown by the darker-colored points, close passes down to less than 1 km occur regularly over the course of the day, especially at latitudes near the "top" or turnaround point (many orbits in this constellation have inclination between 50° and 60°).

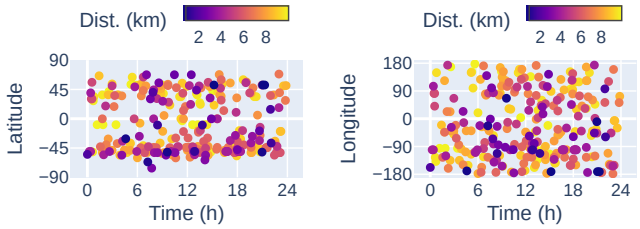


Figure 4: Location of instances where two satellites pass within 10 km of each other over a 24-hour simulation of a TinyLEO constellation.

Although LSN satellites may be small and each particular close encounter may represent a relatively low collision probability, because uncontrolled close encounters are *built into* TinyLEO's geometry, the cumulative probability of collision over a full constellation deployment lifetime approaches 100%. Moreover, collision avoidance maneuvers would have to be computed on-line and on an ad-hoc basis for each close encounter, reducing the constellation's ability to perform avoidance maneuvers in the face of close encounters with *external* objects. As such, the ability of TinyLEO to reduce the number of satellites required compared to other Walker-delta alternatives represents an illusory result that is not feasibly attainable in practice.

3 Principles for LSN Research

In order to address the emergent set of research problems (*e.g.,* minimizing environmental and economic costs while meeting networking requirements), future LSN research must inevitably explore possibilities of new satellite constellations that do not yet exist. As illustrated by the examples above, doing so will require the community to build new intuitions about how space-based infrastructure operates. As

an initial step towards building the new intuitions—*i.e.,* the "knowledge hub"—we collect a set of organizing principles that LSN researchers must be aware of as we go forward. Though not all LSN research will require the same level of engagement with each principle, we contend that for most efforts, adherence to one or more of these principles will be critical to ensure future research creativity does not fall into physically-impossible regions of the design space. Ultimately, we hope this approach will allow researchers to dream "correctly" so that research effort can eventually translate into real-world impact.

3.1 Principle 1: Avoiding Intra-Constellation Collisions

Definition. Because propagation of satellite positions are approximate, avoiding collisions requires ensuring that the approximate locations of any pair of satellites (in the constellation) always exceeds a minimum "collision avoidance" threshold (see § 2.3). Note that this is a design-time consideration that is separate from the runtime consideration of avoiding collisions with *external* objects (*e.g.,* rocket debris, satellites from other constellations, etc.). Constellation designs that violate this principle will incur disproportionately high runtime overheads to avoid internal collisions, as well as the naturally occurring potential for external collisions. Failure to keep collision rates below a certain threshold is known to lead to catastrophic consequences that may limit humanity's future access to space [29].

Takeaways. LSN proposals must prove that their proposed orbits do not collide with each other at any point in time. Note that this principle has been automatically realized in many prior efforts due to use of well-studied orbital geometries like Walker-delta, but does not hold by default in recent proposals that target novel networking-specific geometries.

3.2 Principle 2: Near-Earth Hazards

Definition. Near-Earth space is not completely empty, but filled with several classes of charged particles that get trapped by the Earth's magnetic field [24]. These particles form distinct fields in 3-dimensional space around the Earth (a.k.a. Van Allen radiation belts [45, 49]) that vary in location and density over time as the Earth's magnetic field interacts with the Solar wind. When satellites pass through these fields of charged particles, their payloads are exposed to ionizing radiation that can cause temporary failure or permanent damage to sensitive electronic components.

Takeaways. The key takeaway for LSN research is to acknowledge that different orbits have different baseline cost profiles. Orbits that encounter more radiation over their lifetimes require more expensive satellite payloads with extra shielding (*e.g.,* carbon fiber, aluminum [34, 46]). Orbits that encounter less radiation can get by with cheaper payloads with less shielding.

3.3 Principle 3: Satellite Power Dynamics

Definition. Satellites are typically equipped with solar panels to power their on-board electronic components (*e.g.,* radio transceivers, network processors, etc.) as well as batteries to smooth over interruptions in solar power supply. Batteries on LEO satellites may be exposed to extremely challenging power dynamics over the course of their lifetimes, charging for ∼45 minutes when the satellites is in Sun light and discharging for ∼45 minutes when the satellite is behind the Earth (*e.g.,* consider that Li-ion batteries typically used in satellites only last $O(1\text{k})$ charge cycles [62]).

Takeaways. LSNs should be careful about the power demands of proposed networking processing and other computational functions. Importantly, proposals seeking to offload significant computation to satellite payloads (*e.g.,* CDN and/or AI workloads) must carefully consider the spatiotemporal dynamics of their power requirements.

3.4 Principle 4: ISL Constraints

Definition. A significant theme in LSN research has been how to leverage the potential of inter-satellite links (ISLs) as an alternative to terrestrial wide-area network infrastructure. ISLs escape common vulnerabilities of terrestrial fiber-optic infrastructure (*e.g.,* fiber cuts, Earthquakes) and, in the case of free-space optic ISLs, potentially enable faster-than-fiber connectivity over long-distance links. However, effectively leveraging ISLs requires careful consideration of per-satellite cost tradeoffs (as discussed in § 2.1), link feasibility (based on line-of-sight between satellites), and link setup latency (potentially up to 30 seconds [9]).

Takeaways. The safest approach is to assume each satellite can only form a small number of links (*e.g.,* three links) and to assume that ISLs are static and do not change as satellites proceed through their orbits. Any proposal that seeks to leverage complex topologies with larger numbers of links per satellite or dynamic ISLs to take advantage of changes in relative position between satellites (*e.g.,* at different latitudes) must carefully account for ISL cost tradeoffs, visibility, and setup time latency constraints.

3.5 Principle 5: Incremental Deployment

Definition. Unlike terrestrial network infrastructure, the deployment stage of LSNs faces complex costs and constraints. In particular, satellites must be be launched atop booster vehicles and maneuvered into the desired orbital trajectories. This implies LSN operators must carefully decide when and

from where to schedule launches and how to maneuver satellites from the initial launch orbit into their specified orbits in the constellation geometry.

Takeaways. Deployment overheads should be included in consideration of aggregate LSN costs. For example, certain constellations may be less expensive to deploy while achieving slightly lower networking performance and hence achieve a better overall cost-benefit tradeoff compared to constellations that are more expensive to deploy while achieving better networking performance.

3.6 Principle 6: Orbital Dynamics

Definition. In this context, orbital dynamics refers to the systems of equations that determine where particular orbital bodies, in particular satellites, will be at all times. As discussed in § 2.2, reasoning about how orbital dynamics impact satellite mobility from the perspective of distinct points on the Earth surface requires understanding complex multibody dynamics in three-dimensional space. When an LSN proposal fails to correctly account for orbital dynamics, any assumptions made about where satellites will be at particular points in time may be invalid.

Takeaways. LSN researchers must develop a basic understanding of common orbital dynamic concepts including the intimate relationship between space and time and specific forces that impact both short and long term evolution of orbital parameters (*e.g.,* nodal regression).

4 Discussion

This paper represents an initial step toward a broader reorientation of LSN research. In particular, it places networking innovations in the context of the physical and operational constraints of space-based infrastructure. The principles presented in § 3 reflect our current understanding and experience, but they are neither exhaustive nor intended to be definitive. As satellite technologies evolve and new applications emerge, three aspects need to be evolved: (1) additional constraints will become relevant, (2) existing assumptions will need to be revisited, and (3) some of the principles proposed here may require refinement.

To support this continuing effort, we have established SpaceNet Hub (https://spacenet-hub.github.io/). Our vision is to bootstrap an open-source set of resources and community-of-practice for space networking research. Our hope is that the website serves as a living knowledge hub that organizes the principles introduced in this paper (and additional principles that we may not be aware of yet), provides supporting references and resources, and invites contributions from researchers across networking, aerospace engineering, and related disciplines. Community members can propose new principles, challenge or refine existing ones, contribute relevant literature, and discuss how space-specific constraints should inform networking research. By maintaining these resources openly, we aim to make knowledge from outside the networking community more accessible and encourage researchers to critically examine the assumptions underlying their system designs.

We envision SpaceNet Hub as a resource that grows alongside the field, rather than a fixed collection of design rules. For example, growing interest in deploying CDNs and data centers in space raises questions about how orbital altitude, satellite power dynamics (Principle 3), communication latency, and network connectivity jointly constrain infrastructure placement and operation. Addressing such questions requires more than optimizing individual networking metrics; it requires understanding how networking decisions interact with spacecraft design, mission feasibility, and business objectives. We hope the principles developed here, together with their continued refinement through SpaceNet Hub, will help researchers formulate more realistic problems, evaluate proposed solutions against near-Earth space constraints, and identify new opportunities at the intersection of networking and space systems.

Acknowledgments

We thank our shepherd Jon Crowcroft and the anonymous reviewers for their constructive feedback. This work is supported by the National Science Foundation (CNS-2145813). 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 NSF.

References

[1] [n. d.]. CelesTrak. https://celestrak.org/. Accessed: 2024-12.
[2] [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.
[3] [n. d.]. Starlink. https://starlink.com/. Accessed 2026-05.
[4] [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.
[5] Paul Adams. [n. d.]. Musk cuts Starlink access for Russian forces, giving Ukraine an edge at the front. https://www.bbc.com/news/articles/c0q3ndj7052o. Accessed: 2026-04.
[6] Ian F Akyildiz, Eylem Ekici, and Michael D Bender. 2002. MLSR: A novel routing algorithm for multilayered satellite IP networks. *IEEE/ACM Transactions on networking* 10, 3 (2002), 411–424.
[7] 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.
[8] Yang Cao, Shao-Yu Lien, and Ying-Chang Liang. 2021. Multi-tier collaborative deep reinforcement learning for non-terrestrial network

empowered vehicular connections. In *2021 IEEE 29th International Conference on Network Protocols (ICNP)*. IEEE, 1–6.

[9] Carlos Carrizo, Markus Knapek, Joachim Horwath, Dionisio Gonzalez, and Paul Cornwell. 2020. Optical inter-satellite link terminals for next generation satellite constellations. 2. doi:10.1117/12.2545629

[10] Andrew Cavalier. 2026. Why Orbital Data Centers are so Hard: The Idea Might Run Aground on the Second Law of Thermodynamics. *IEEE Spectrum* 63, 7 (2026), 18–23.

[11] Aizaz U Chaudhry and Halim Yanikomeroglu. 2021. Laser intersatellite links in a starlink constellation: A classification and analysis. *IEEE vehicular technology magazine* 16, 2 (2021), 48–56.

[12] Aizaz U Chaudhry and Halim Yanikomeroglu. 2022. Optical wireless satellite networks versus optical fiber terrestrial networks: The latency perspective: Invited chapter. In *30th Biennial Symposium on Communications 2021*. Springer, 225–234.

[13] Quan Chen, Giovanni Giambene, Lei Yang, Chengguang Fan, and Xiaoqian Chen. 2021. Analysis of inter-satellite link paths for LEO mega-constellation networks. *IEEE Transactions on Vehicular Technology* 70, 3 (2021), 2743–2755.

[14] Qian Chen, Weixiao Meng, Tony QS Quek, and Shuyi Chen. 2022. Multi-tier hybrid offloading for computation-aware IoT applications in civil aircraft-augmented SAGIN. *IEEE Journal on Selected Areas in Communications* 41, 2 (2022), 399–417.

[15] Stephen Chen. [n. d.]. Starlink breaks space traffic safety rules and China may follow: study. https://www.scmp.com/news/china/science/article/3201365/starlink-breaks-space-traffic-safety-rules-and-china-may-follow-study. Accessed: 2026-05.

[16] 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.

[17] 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.

[18] Kenneth E Crowe. 1999. *A comparative analysis of the iridium and globalstar satellite transmission paths*. Technical Report.

[19] Cui-Qin Dai, Guangyan Liao, and Qianbin Chen. 2021. Service-oriented routing with Markov space-time graph in low earth orbit satellite networks. *Transactions on Emerging Telecommunications Technologies* 32, 7 (2021), e4072.

[20] Ruoqi Deng, Boya Di, Hongliang Zhang, Linling Kuang, and Lingyang Song. 2021. Ultra-dense LEO satellite constellations: How many LEO satellites do we need? *IEEE Transactions on wireless communications* 20, 8 (2021), 4843–4857.

[21] FCC. [n. d.]. Satellite communications services (FCC Report No. SES-02767). https://docs.fcc.gov/public/attachments/DOC-411304A1.pdf. Accessed: 2025-07.

[22] 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.

[23] Xiangqiang Gao, Rongke Liu, and Aryan Kaushik. 2022. Virtual network function placement in satellite edge computing with a potential game approach. *IEEE Transactions on Network and Service Management* 19, 2 (2022), 1243–1259.

[24] Timothy B Guild, Scott C Davis, Alex J Boyd, T Paul O'Brien, Joseph F Fennell, Joseph E Mazur, Douglas G Brinkman, Glenn E Peterson, and Alan B Jenkin. 2022. *Best Practices for Generating Space Environment Specification with Modern Tools*. Aerospace Corporation.

[25] 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.

[26] 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*. 205–211.

[27] Rachel Jewett. 2021. SpaceX Resumes Starlink Launches with Optical Link-Equipped Satellites. https://www.satellitetoday.com/launch/2021/09/14/spacex-resumes-starlink-launches-with-optical-link-equipped-satellites/. Accessed: 2026-05.

[28] 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.

[29] Donald J Kessler and Burton G Cour-Palais. 1978. Collision frequency of artificial satellites: The creation of a debris belt. *Journal of Geophysical Research: Space Physics* 83, A6 (1978), 2637–2646.

[30] 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.

[31] 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).

[32] 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. 521-530 pages. doi:10.1109/INFOCOM52122.2024.10621083

[33] 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.

[34] Matthew Large. [n. d.]. European Space Radiation Shielding Workshop. https://indico.esa.int/event/501/contributions/9646/. Accessed: 2026-05.

[35] Israel Leyva-Mayorga, Beatriz Soret, and Petar Popovski. 2021. Inter-plane inter-satellite connectivity in dense LEO constellations. *IEEE Transactions on Wireless Communications* 20, 6 (2021), 3430–3443.

[36] 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.

[37] Yuanjie Li, Yimei Chen, Jiabo Yang, Jinyao Zhang, Bowen Sun, Lixin Liu, Hewu Li, Jianping Wu, Zeqi Lai, Qian Wu, et al. 2025. Small-scale LEO Satellite Networking for Global-scale Demands. In *Proceedings of the ACM SIGCOMM 2025 Conference*. 917–931.

[38] 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.

[39] Miles Lifson, David Arnas, Martín Avendaño, and Richard Linares. 2025. Low-Earth-Orbit Packing: Implications for Orbit Design and Policy. *Journal of Spacecraft and Rockets* 62, 2 (2025), 452–465.

[40] Xingqin Lin, Stefan Rommer, Sebastian Euler, Emre A Yavuz, and Robert S Karlsson. 2021. 5G from space: An overview of 3GPP non-terrestrial networks. *IEEE Communications Standards Magazine* 5, 4 (2021), 147–153.

[41] 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.

[42] Wei Liu, Yuanjie Li, Jingyi Lan, Hewu Li, Yimei Chen, Lixin Liu, Jiabo Yang, Xi Long, Li Ouyang, Minghao Tang, et al. 2025. Direct-to-Cell Satellite Network without Satellite Navigation. In *Proceedings of the ACM SIGCOMM 2025 Conference*. 932–947.

[43] Yu Liu, Luhan Wang, Zhaoming Lu, Keliang Du, and Guochu Shou. 2023. A stateless design of satellite–terrestrial integrated core network and its deployment strategy. *IEEE Transactions on Network and Service Management* 21, 1 (2023), 953–966.

[44] Doug Madory. [n. d.]. Starlink Enters Transit Market With Community Gateways. https://www.kentik.com/blog/starlink-enters-transit-market-with-community-gateways/. Accessed: 2026-04.

[45] 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.

[46] Masood Mayanbari and Yaser Kasesaz. 2011. Design and analyse space radiation shielding for a nanosatellite in Low Earth Orbit (LEO). In *Proceedings of 5th International Conference on Recent Advances in Space Technologies-RAST2011*. IEEE, 489–493.

[47] 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.

[48] 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.

[49] Chris Misa and Ramakrishnan Durairajan. 2025. The Internet from Space, Reimagined: Leveraging Altitude for Efficient Global Coverage. In *Proceedings of the 3rd Workshop on LEO Networking and Communication*. 57–63.

[50] Chris Misa and Ramakrishnan Durairajan. 2025. Sustainability or Survivability? Eliminating the Need to Choose in LEO Satellite Constellations. In *Proceedings of the 24th ACM Workshop on Hot Topics in Networks*. 123–130.

[51] Nitinder Mohan, Andrew E Ferguson, Hendrik Cech, Rohan Bose, Prakita Rayyan Renatin, Mahesh K Marina, and Jörg Ott. 2024. A multifaceted look at starlink performance. In *Proceedings of the ACM Web Conference 2024*. 2723–2734.

[52] Robert Muggah. [n. d.]. Starlink Has Privatized Geopolitics. https://foreignpolicy.com/2026/03/20/starlink-spacex-musk-geopolitics-war-ukraine-russia-iran/. Accessed: 2026-04.

[53] Robin Ohs, Gregory F Stock, Andreas Schmidt, Juan A Fraire, Jörg Ott, and Holger Hermanns. 2026. Dark Clouds Rising in Low-Earth Orbit: On Environmental Limits to Massive Orbital AI. In *Proceedings of the ACM SIGCOMM 2026 Conference*. 1960–1965.

[54] 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.

[55] Prashant Raman, Duncan Stewart, Gillian Crossan, Tim Bottke, Jody McDermott, and Ben Stanton. [n. d.]. Next-gen satellite Internet is transforming pricing, capacity, and regulation worldwide. https://www.deloitte.com/us/en/insights/industry/technology/technology-media-and-telecom-predictions/2026/next-gen-satellite-internet.html. Accessed: 2026-04.

[56] Thomas Royster, Jun Sun, Aradhana Narula-Tam, and Thomas Shake. 2023. Network performance of pLEO topologies in a high-inclination Walker delta satellite constellation. In *2023 IEEE Aerospace Conference*. IEEE, 1–9.

[57] Vaibhav Singh, Tusher Chakraborty, Suraj Jog, Om Chabra, Deepak Vasisht, and Ranveer Chandra. 2024. Spectrumize: Spectrum-efficient satellite networks for the internet of things. In *21st USENIX Symposium on Networked Systems Design and Implementation (NSDI 24)*. 825–840.

[58] Starlink. 2026. Progress 2024. https://starlink.com/public-files/starlinkProgressReport_2024.pdf. Accessed: 2026-04.

[59] Starlink. 2026. Technology. https://web.archive.org/web/20260321220627/https://starlink.com/technology. Accessed: 2026-04.

[60] 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.

[61] Hammas Bin Tanveer, Mike Puchol, Rachee Singh, Antonio Bianchi, and Rishab Nithyanand. 2023. Making sense of constellations: Methodologies for understanding starlink's scheduling algorithms. In *Companion of the 19th International Conference on Emerging Networking EXperiments and Technologies*. 37–43.

[62] Battery University. [n. d.]. BU-808: How to Prolong Lithium-based Batteries. https://www.batteryuniversity.com/article/bu-808-how-to-prolong-lithium-based-batteries/. Accessed: 2026-05.

[63] David A Vallado. 2001. *Fundamentals of astrodynamics and applications*. Vol. 12. Springer Science & Business Media.

[64] John G Walker. 1982. *Coverage Predictions and Selection Criteria for Satellite Constellations*. Technical Report.

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

[66] Markus Werner, Cecilia Delucchi, H-J Vogel, Gerard Maral, and J-J De Ridder. 1997. ATM-based routing in LEO/MEO satellite networks with intersatellite links. *IEEE Journal on Selected areas in Communications* 15, 1 (1997), 69–82.

[67] Hao Wu, Yizhan Han, Mohit Rajpal, Qizhen Zhang, and Jingxian Wang. 2025. SaTE: Low-Latency Traffic Engineering for Satellite Networks. In *Proceedings of the ACM SIGCOMM 2025 Conference*. 896–916.

[68] Junyi Yang, Bo Li, Kexin Fan, Lirong An, and Qinyu Zhang. 2024. Analysis of laser intersatellite links and topology design for mega-constellation networks. *IEEE internet of things journal* 11, 21 (2024), 34554–34566.

[69] Jinwei Zhao and Jianping Pan. 2024. LENS: A LEO satellite network measurement dataset. In *Proceedings of the 15th ACM Multimedia Systems Conference*. 278–284.

[70] William X Zheng, Aryan Taneja, Maleeha Masood, Anirudh Sabnis, Ramesh K Sitaraman, and Deepak Vasisht. 2025. StarCDN: Moving Content Delivery Networks to Space. In *Proceedings of the ACM SIGCOMM 2025 Conference*. 948–962.

[71] Rachael Zisk. 2022. Musk Gives Starlink Update. https://payloadspace.com/musk-gives-starlink-update/. Accessed: 2026-05.