Rethinking LEO Networks for Users in Orbit

Jiabo Yang*, Lixin Liu*, Yuanjie Li, Hewu Li, Qian Wu, Zeqi Lai, Jun Liu
Tsinghua University, Zhongguancun Laboratory

Abstract

Satellites are becoming not only network infrastructure but also network users. The rise of Earth observation, onboard artificial intelligence, and orbital data centers is creating substantial demand for space data backhaul. Traditional backhaul through direct contacts with ground stations is unlikely to support this growth because it suffers from long waiting delays, limited contact windows, and bandwidth constraints. We argue for reusing existing LEO satellite networks as an orbital backhaul alternative. Their broad coverage, low latency, and high capacity make them a promising substrate for delivering data from orbit to the ground. However, satellites as space users differ substantially from terrestrial users in demand distribution, access dynamics, and forwarding requirements. These differences challenge the assumptions behind existing LEO satellite network designs. This paper systematically examines these challenges and aims to motivate future research on reusing LEO satellite networks to serve users in orbit.

CCS Concepts

• **Networks** → **Network architectures**.

Keywords

LEO satellite networks, orbital backhaul

ACM Reference Format:
Jiabo Yang, Lixin Liu, Yuanjie Li, Hewu Li, Qian Wu, Zeqi Lai, Jun Liu. 2026. Rethinking LEO Networks for Users in Orbit. 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, 8 pages. https://doi.org/10.1145/3847290.3859933

*The first two authors contribute equally to this work. Corresponding author: Yuanjie Li (yuanjiel@tsinghua.edu.cn).



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/10.1145/3847290.3859933

1 Introduction

Satellites are becoming not only part of the network infrastructure but also network users. Traditional Earth observation satellites already generate large volumes of sensing data, while advances in onboard computing are enabling more complex processing and artificial intelligence workloads in orbit. SpaceX has proposed Starmind, a constellation of up to one million satellites [14, 22]. Starcloud and Blue Origin Sunrise have outlined plans for systems involving 88,000 and 51,600 satellites, respectively [1, 7, 19].

These emerging space users create a growing demand for backhaul with high capacity and low latency. They generate and process data driven by their missions. They must eventually deliver the resulting data to the ground for terrestrial applications. The volume can already reach tens to hundreds of terabytes per day: ESA reports that Earth observation satellites collectively generate more than 150 TB of data per day, while NASA estimates that the NISAR mission alone will produce about 80 TB of data products per day [5, 12].

Today's space users primarily rely on direct ground station contacts for data backhaul. The data is then forwarded over the terrestrial Internet to its final destination. However, each contact is limited by orbital visibility, antenna schedules, and the distribution of ground stations. A satellite may generate data continuously but wait for a suitable ground station before transmission. This waiting time can dominate end-to-end latency, making direct downlink poorly suited to future applications that require timely access to orbital data. Distributed ground station networks and onboard data reduction can improve delivery opportunities, but they remain constrained by intermittent access to the ground. Future space users therefore need a backhaul service that can carry large data volumes without relying solely on direct ground station contacts.

A promising alternative is to use LEO satellite networks to provide orbital backhaul. As illustrated in Figure 1, a space user could first transmit data to a nearby LEO network satellite. The network could then forward the traffic through its inter-satellite links (ISLs) to a ground gateway with available capacity. The traffic is then forwarded over the terrestrial Internet to its final destination. Unlike direct downlink, this approach allows orbital users to backhaul data via gateways beyond their direct visibility.

We argue that existing LEO satellite networks can be reused for orbital backhaul. Current networks are already

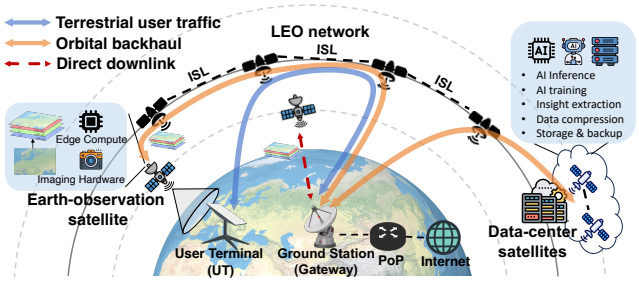


Figure 1: Reusing existing LEO satellite networks for orbital backhaul.

equipped with radio frequency (RF) and optical interfaces, inter-satellite connectivity, and globally distributed ground gateways. Their terrestrial access demand is also highly uneven. Satellites over oceans and sparsely populated regions can retain access capacity that is not consumed by terrestrial users. A comparison of representative SSO Earth observation demand and reported interface capacity suggests that raw access capacity may not be the primary limitation. Reuse also avoids deploying another global constellation solely for orbital backhaul. It preserves increasingly scarce orbital resources, avoids duplicating space and ground infrastructure, and can provide service sooner than a dedicated network that requires new authorization, manufacturing, launch, and deployment. Orbital backhaul could thus become an additional service of LEO satellite networks, extending their user population from the Earth to space.

While promising, reusing existing LEO satellite networks for orbital backhaul is not straightforward. These networks are designed primarily for terrestrial users. Orbital users differ from terrestrial users in their mobility patterns and traffic demands. Supporting both types of users on the same network challenges three existing design assumptions. First, terrestrial access resources are provisioned around population-based demand, whereas orbital demand follows mission-driven orbits and changes with relative satellite motion. Limited radio and optical interfaces must also be shared among terrestrial users, orbital users, and communication satellites. Second, terrestrial users provide a stable geographic reference, and ISLs are commonly formed between satellites with similar orbits. Links to space users on heterogeneous orbits exhibit high relative motion, short contacts, and frequent handovers. Third, terrestrial users can rely on relatively stable gateways for data delivery, whereas the preferred gateways for orbital users change continuously with their motion. Dynamic gateway selection must also be coordinated with routing across the LEO network, without degrading terrestrial services.

Through empirical analysis of SSO Earth observation satellites and an operational LEO network, we quantify these mismatches and identify key challenges in access provisioning, link management, and data delivery. We hope these findings will motivate further research on extending existing LEO networks to serve users in orbit.

2 The Rise of Users in Orbit

This section characterizes the emerging users in orbit (§2.1) and the backhaul demand they create (§2.2).

Table 1: Proposed orbital parameters of representative orbital data-center constellations.

System	Group	Satellites	H (km)	i	Orbit type
SpaceX [14]	1	99,900	550–568	26–32°	Low inclination
	2	99,980	565–585	97.7°	Dawn-dusk SSO
	3	249,750	686–718	30°	Low inclination
	4	244,860	707–744	97.2°	Dawn-dusk SSO
	5	249,750	946–978	30°	Low inclination
	6	253,880	967–1,002	99.4°	Dawn-dusk SSO
Starcloud [7, 19]	–	88,000	600–850	97.8–98.8°†	Dawn-dusk SSO
Blue Origin Sunrise [1]	–	51,600	500–1,800	97–104°	SSO

†Estimated from the disclosed SSO altitude range.

2.1 Users in Mission-Driven Orbits

Users in orbit are growing from a small number of specialized mission satellites into constellation-scale populations of data-intensive spacecraft. Early space users were mainly Earth observation satellites. A limited number of satellites collected imagery and sensor measurements for applications such as environmental monitoring, weather forecasting, disaster response, and mapping. Recent advances in onboard computing and artificial intelligence have expanded these missions to include data filtering, inference, and analysis in orbit. More recent proposals go further by deploying large orbital computing and data center constellations (Table 1). SpaceX has proposed Starmind, which could contain up to one million satellites [22]. Starcloud and Blue Origin Sunrise have disclosed systems involving 88,000 and 51,600 satellites, respectively [1, 7, 19]. These proposals broaden the role of space users from sensing to persistent computation. Users in orbit are therefore becoming a constellation-scale population rather than a collection of isolated spacecraft.

These satellites are mission-driven, so their orbital parameters are not chosen arbitrarily. Factors such as geographic coverage, revisit frequency, lighting conditions, and solar energy determine where satellites operate. Sun-synchronous orbit (SSO) is particularly attractive because its orbital plane maintains a nearly constant orientation relative to the Sun, as illustrated in Figure 2. A satellite therefore passes over the same region at approximately the same local solar time and observes it under consistent lighting conditions, which is

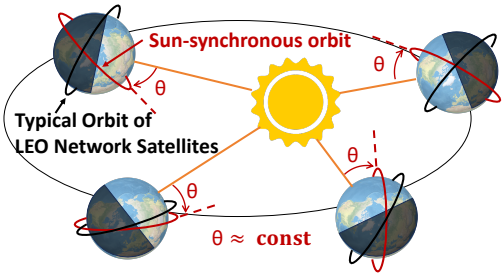
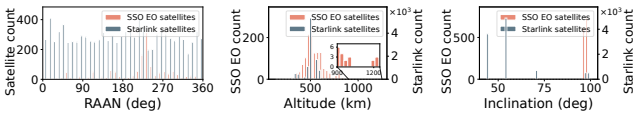


Figure 2: Sun-synchronous orbit.



(a) RAAN **(b) Altitude** **(c) Inclination**

Figure 3: Orbital heterogeneity between SSO space users and LEO network satellites.

valuable for Earth observation [6]. Dawn-dusk SSO further keeps the satellite close to the day-night boundary, providing near-continuous solar exposure, which benefits computing systems with high energy demand. Accordingly, SpaceX and Starcloud plan to place large portions of their computing constellations in such orbits [7, 14, 19].

Figure 3 reveals orbital heterogeneity between SSO Earth observation satellites and operational Starlink satellites. SSO satellites exhibit a highly non-uniform RAAN distribution, span a broader altitude range, and concentrate near inclinations of 97–99 degrees. Starlink satellites instead distribute at several discrete altitudes and inclination groups. Despite partial overlap, the two types of satellites follow substantially different orbital distributions.

2.2 Growing Demand for Backhaul

The growing population of space users creates substantial demand for data backhaul. Earth observation is a typical example. As illustrated in Figure 4, an Earth observation satellite collects and processes imagery or sensor data, stores the results onboard, and later transfers them to terrestrial computing facilities. Existing Earth observation systems can generate hundreds of gigabytes to several terabytes per satellite each day [4, 20, 21]. A workload of 1 TB per day corresponds to an average rate of approximately 92.6 Mbps. We estimate that, across the 1,363 SSO Earth observation satellites from Skyrocket [9], this yields an aggregate average demand of 126.2 Gbps. The required rate during short contact windows would be considerably higher. Future orbital AI and data center constellations will further increase this

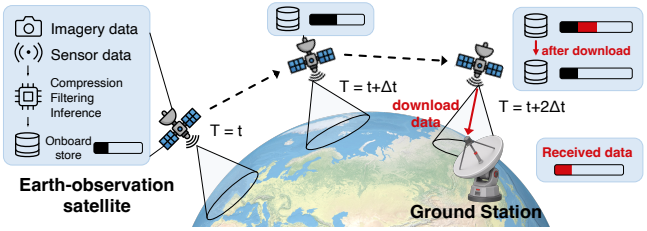


Figure 4: The workflow of Earth-observation satellites.

demand. They may need to exchange model parameters, intermediate results, selected raw data, and operational status with ground infrastructure.

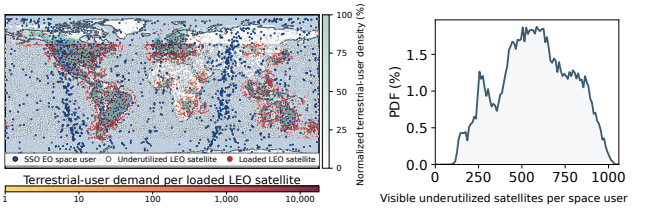
However, existing space missions still rely primarily on direct ground station contacts for data backhaul. A satellite must wait until a suitable station becomes visible and then transfer data within a limited contact window. This store-and-forward process introduces long delays, bursty transmission, and competition for antennas, spectrum, and schedules. Adding ground stations is costly and constrained by geography and regulation. Systems such as L2D2 [21] and Serval [20] improve the use of existing infrastructure, but remain limited by intermittent ground contact. As users in orbit continue to grow, orbital backhaul requires a communication substrate that provides broader availability and greater performance without relying solely on direct contacts with ground stations.

3 Reusing LEO Networks for Orbital Backhaul

We argue that existing LEO satellite networks are a promising substrate for orbital backhaul. Their interfaces and residual capacity make reuse feasible (§3.1), while avoiding the cost and delay of dedicated backhaul constellations (§3.2).

3.1 Existing LEO Networks Can Carry More

Existing interfaces enable physical access. Today's LEO satellites are equipped with radio frequency (RF) and optical interfaces that could be adapted to serve users in orbit. Earth observation satellites commonly use the S, X, and Ka bands for telemetry, control, and data downlink [8], while broadband LEO networks operate in bands including Ku and Ka [16]. These overlapping capabilities provide a potential path for radio access between the two systems. For example, SpaceX filings describe each Starlink V2 satellite as supporting 24 independently steerable radio beams and three optical links with a capacity of 200 Gbps each [16]. The upcoming Starlink V3 satellites further expand this capability by increasing the number of optical intersatellite links from three to six per satellite [17]. SpaceX also offers the Mini Laser for integration with third-party spacecraft, with



(a) Terrestrial access demand (b) Visible underutilized sats

Figure 5: Uneven terrestrial demand and underutilized LEO satellites.

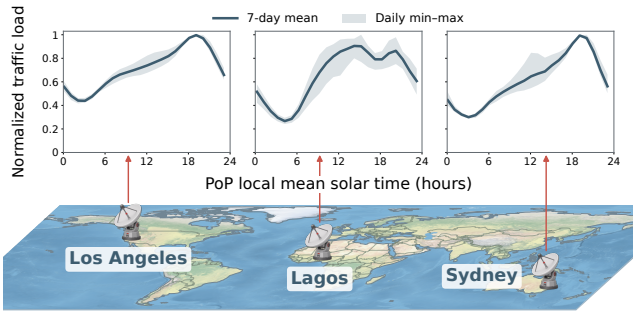


Figure 6: Diurnal traffic variation at LEO network PoPs.

a reported rate of up to 25 Gbps [13, 18]. Once a space user connects to the network, existing ISLs and ground gateways can carry its traffic to terrestrial infrastructure. These interfaces provide the basic access and forwarding capabilities required for orbital backhaul.

Reported interface capacity suggests potential. Using the demand estimate in §2.2, the 1,363 SSO Earth observation satellites require an aggregate average rate of 126.2 Gbps. The average demand per satellite, 92.6 Mbps, is far below the reported 25 Gbps rate of a Mini Laser terminal. For RF access, 57 satellites with 24 beams each would offer 1,368 beams simultaneously, comparable to the current space user population. These comparisons suggest that existing interfaces have the potential to support orbital backhaul. However, their actual capacity depends on interface availability, resource sharing, and network conditions. Continued advances in RF and optical interfaces, such as Starlink V3's substantially higher radio capacity and six optical ISLs per satellite, suggest that LEO networks could further support orbital backhaul at larger scales [17].

Spatially uneven demand creates reuse opportunities. Terrestrial demand is unevenly distributed across geographic regions [2, 10, 11]. Figure 5a shows that satellites over densely populated regions carry more traffic, while many satellites over oceans and sparsely populated regions remain lightly loaded. Prior work shows that 90% of satellite capacity remains underutilized [11]. Meanwhile, orbital users move continuously across regions with different terrestrial loads. Figure 5b reports the number of geometrically visible, lightly loaded LEO satellites for each SSO Earth observation satellite. It shows that 96.9% of these space users can see at least 200 lightly loaded LEO satellites. These results suggest opportunities to reuse underutilized access resources for orbital backhaul. With inter-satellite connectivity, orbital traffic could reach ground gateways beyond the user's direct visibility.

Temporal variation creates reuse opportunities. Terrestrial traffic usually varies over time, often following daily activity patterns. This variation creates opportunities to reuse network resources during periods of lower terrestrial demand. Figure 6 shows estimated traffic profiles of three representative Starlink points of presence (PoPs) based on Cloudflare measurements [3]. All three exhibit substantial diurnal variation. As these PoPs span different time zones, their periods of lower demand occur at different times. These results suggest opportunities to deliver orbital traffic through less-loaded ground infrastructure. By using ISLs to reach suitable gateways, orbital users could benefit from these variations rather than relying only on direct ground station contacts.

3.2 Why Reuse Existing LEO Networks

We next examine the broader rationale for reusing existing LEO satellite networks for orbital backhaul.

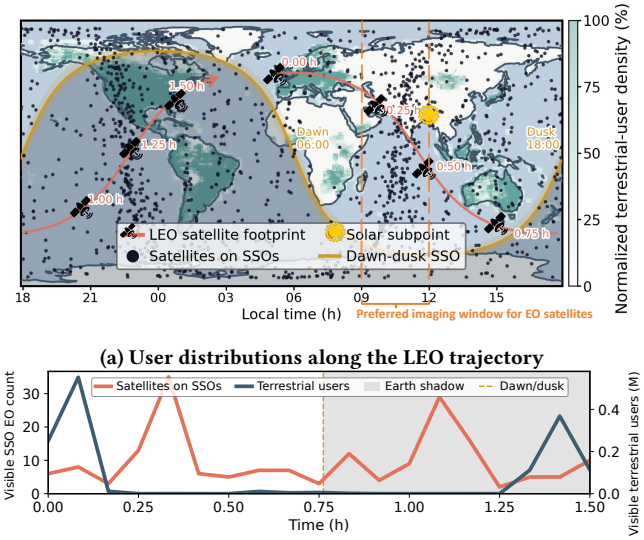
Preserves scarce orbital resources. LEO space is becoming increasingly crowded, including approximately 15,000 active satellites and more than 9,910 debris objects [15]. A dedicated backhaul constellation would add another large population of spacecraft and increase the need for orbit coordination, collision avoidance, and spectrum sharing. Reuse supports emerging orbital users without requiring another global satellite constellation for backhaul.

Avoids high deployment costs. Deploying a dedicated backhaul constellation is costly. Existing LEO satellite networks already provide much of the infrastructure needed for orbital traffic delivery. Reusing this infrastructure can reduce deployment costs.

Shortens time to service. Building a new LEO network requires regulatory approval, spectrum coordination, manufacturing, launch, and ground deployment. These steps can span several years before the network provides broad and stable service. In contrast, reusing existing LEO satellite networks for orbital backhaul can fastly enable emerging orbital applications without waiting for a long time.

4 Challenges in Network Reuse

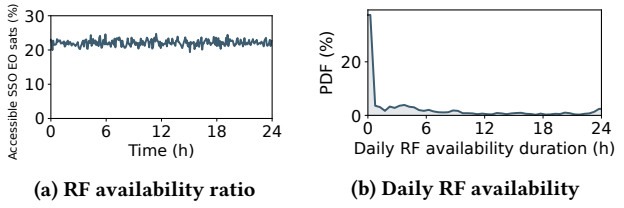
While promising, reusing existing LEO satellite networks to serve orbital users introduces new challenges. We first examine how mission-driven orbital demand complicates access



(a) User distributions along the LEO trajectory

(b) The number of user demand per satellite

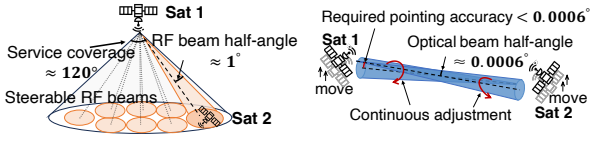
Figure 7: Different access demand patterns of terrestrial and orbital users.



(a) RF availability ratio

(b) Daily RF availability

Figure 8: Mismatch between ground-facing LEO resources and nonuniform SSO space user demand.



(a) RF coverage

(b) Optical pointing

Figure 9: RF coverage and optical pointing constraints.

provisioning (§4.1). We then analyze how heterogeneous orbits challenge link continuity (§4.2). Finally, we show that moving users in orbit makes fixed ground gateway-based routing inefficient for data delivery (§4.3).

4.1 Demand and Provisioning Mismatch

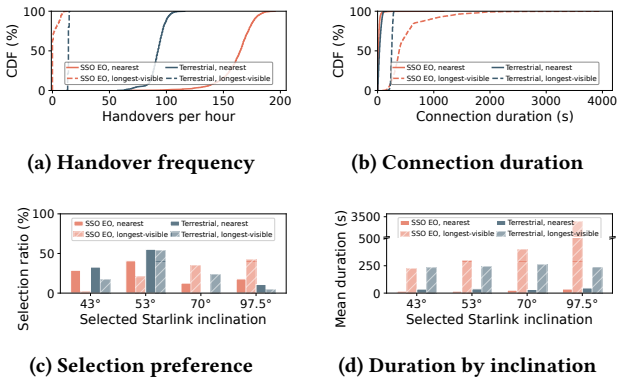
Terrestrial demand follows population distribution. Its major demand regions are geographically stable and concentrated at low and middle latitudes. Existing LEO satellite networks reflect this demand pattern. Satellites are distributed approximately uniformly within each orbital shell to maintain continuous coverage, while many operational satellites occupy low and middle inclination shells that repeatedly cover populated regions. This relatively stable demand geometry allows operators to provision access beams and capacity according to geographic demand.

Space demand follows a different spatial structure. Missions drive satellites to concentrate on orbits with similar altitudes and inclinations. As introduced in §2.1, Earth observation satellites commonly use sun-synchronous orbits to obtain consistent lighting conditions [6]. Recently proposed systems, such as Starmind and Starcloud, also use sun-synchronous orbits to obtain long periods of solar exposure [7]. So space users form concentrated populations along particular mission-driven orbits. Figure 7a shows that this orbital concentration is significantly different from the population-based distribution of terrestrial users.

Orbital demand and access resources are mismatched. This mismatch changes as the LEO network moves. Figure 7b shows that terrestrial and SSO demand vary differently as a LEO satellite moves. Terrestrial demand follows population density, while SSO demand exhibits periodic peaks near regions where orbital users concentrate. A LEO satellite for backhaul may therefore encounter very different terrestrial and orbital demand at the same time. We quantify one consequence of this mismatch under the default ground-facing RF configuration. We define the RF availability ratio as the fraction of SSO users covered by at least one Starlink satellite. As shown in Figure 8a, only about 22% of SSO users are covered on average. We also calculate the total coverage duration of each user over one day. Figure 8b shows that most users are covered for less than six hours. Thus, even before resource contention, the default ground-facing RF configuration leaves many orbital users without coverage.

We next examine that even when orbital users are reachable by adjusting the interface orientation, the available interfaces can serve only a limited number of them at the same time. Figure 9 illustrates the different coverage and pointing constraints of RF and optical access. RF access supports more concurrent users, but the number of available beams is still limited, and each connection receives less capacity. Optical links provide much higher capacity, but each terminal can point to only one user at a time, and the number of terminals on each satellite is limited (3-6 ISLs) [24].

This challenge is further exacerbated because the same limited resources already support terrestrial users and network connectivity. Reusing them for orbital users creates direct competition between terrestrial and orbital demand. Access provisioning must continuously match users in orbit with suitable interfaces and capacity while preserving terrestrial access and inter-satellite connectivity. Coordinating these limited shared resources across two distinct demand patterns poses a significant challenge.



(a) Handover frequency (b) Connection duration

(c) Selection preference (d) Duration by inclination

Figure 10: The comparison of link continuity.

4.2 Heterogeneous Link Dynamics

Terrestrial connectivity has stable references. Existing LEO networks manage terrestrial access using geographic regions. A terrestrial user may move, but its position usually changes much more slowly than the LEO satellite. Changes in its access link arise mainly from satellite motion. Although the serving satellite changes, the user remains within a relatively stable service region with a fixed gateway. This region provides a persistent reference for access satellite selection and handover [25]. In the constellation, inter-satellite links are mainly established between neighboring satellites with similar orbital parameters. Although these satellites keep moving, their relative geometry changes slowly, allowing their links to remain stable for long periods.

Heterogeneous orbits destabilize orbital access. An orbital user and a LEO network satellite may follow different orbits. Their altitude, inclination, and phase can differ substantially. These differences produce high relative motion and short contacts, even though both satellites follow predictable trajectories. Figure 10a compares the handover frequencies of terrestrial and SSO users. SSO users switch their serving satellites much more frequently than terrestrial users when accessing the nearest satellite. Figure 10b compares the resulting link dynamics. Under nearest satellite selection, SSO users experience roughly twice as many handovers as terrestrial users, while most connections last only tens of seconds. The longest visible strategy substantially reduces the handover frequency and extends many SSO connections to hundreds of seconds, with some lasting more than 2,000 seconds. The selection strategy strongly affects how long each link can be maintained.

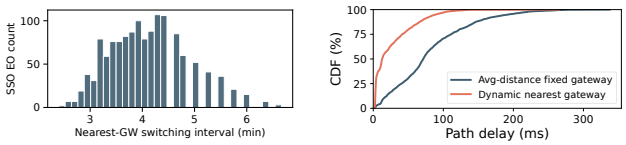
Multiple orbital shells further complicate serving satellite selection. Figure 10c shows that terrestrial users are served mainly by the 43° and 53° shells under both selection strategies. These shells provide repeated coverage over the low and middle latitude regions where terrestrial demand is concentrated. Under nearest satellite selection, SSO users distribute their connections relatively evenly across the four shells. The result changes under longest visible association. More than 42.46% of SSO connections select the 97.5° shell, whose inclination closely matches that of SSO users. Figure 10d explains this preference. Under nearest satellite selection, the mean connection duration remains only tens of seconds across all shells. Under longest visible selection, connections to the 97.5° shell last more than 3,338 seconds on average. This duration is substantially longer than that achieved through the other shells. Orbital similarity, rather than proximity alone, is therefore a key factor in link continuity.

These results expose a fundamental challenge for link management. The closest satellite may provide a strong instantaneous connection but require frequent handovers. A satellite with compatible orbital motion may remain visible much longer, but identifying it requires knowledge of future contact opportunities and resource availability. The candidate set also changes as the user moves across orbital shells. The network must repeatedly select among proximity, contact duration, and available capacity while limiting service interruption and control overhead. For optical access, each change may additionally require acquisition, pointing, and tracking of a new target. Maintaining continuous service across heterogeneous and rapidly changing orbital relationships is therefore substantially more difficult than managing terrestrial access or links within a homogeneous orbital shell.

4.3 Backhaul without Fixed Gateway

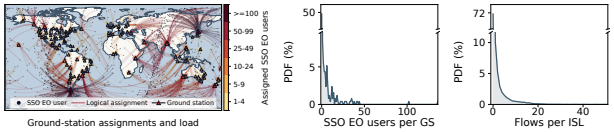
Terrestrial delivery relies on stable gateways. Existing LEO networks organize data delivery according to geographical regions. Although a terrestrial user's serving satellite changes, the user usually remains within the same geographic region. So the network can deliver its traffic through a relatively stable gateway and home PoP. For instance, Starlink satellites forward user traffic to a home PoP before entering the public Internet, and the user's IPv6 prefix is associated with its PoP [23]. This stable assignment simplifies addressing and routing because network state does not need to update every satellite handover.

Orbital motion changes the preferred gateway. A fixed gateway is less suitable for a user in orbit. As the user moves, the best available gateway to receive its traffic changes continuously. Figure 11a shows the interval between changes in the nearest gateway for SSO users. The nearest gateway changes about every four minutes on average. We quantify the cost of retaining a fixed gateway by comparing two delivery strategies. The fixed strategy selects one gateway based on its average distance over the trace and retains it throughout the experiment. The dynamic strategy selects the nearest gateway at each time. As shown in Figure 11b, the fixed strategy produces an average path delay of 81.18



(a) Gateway switching (b) Path delay

Figure 11: Gateway changes and path inflation.



(a) Gateway assignments (b) GS load (c) ISL load

Figure 12: Uneven gateway and ISL loads under shortest-path routing.

ms, compared with 26.77 ms under dynamic selection. Selecting the nearest gateway therefore reduces the average path delay by approximately 67%. These results highlight the benefit of adaptive gateway selection for orbital users. However, changing the gateway or PoP may require prefix reassignment and routing updates, which can disrupt ongoing sessions. So the user's address and session state must remain stable even when the ground gateway changes.

Gateway selection alone is insufficient. Dynamic gateway selection can reduce path delay, but it does not guarantee efficient delivery. Figure 12a maps the gateway assignments under nearest gateway selection, while Figure 12b shows the resulting load distribution. It shows that traffic concentrates on a subset of gateways because users in similar orbits often prefer the same ground locations. These gateways may become congested while others remain lightly loaded. The resulting queueing delay may offset the propagation delay saved by choosing the nearest gateway.

Meanwhile, orbital traffic may traverse multiple inter-satellite links before reaching the selected gateway. Because orbital demand can be both large and concentrated along particular orbits, shortest path routing may funnel traffic onto a small subset of satellites and links. Figure 12c shows that routing SSO traffic along shortest paths to the nearest gateways produces a highly uneven inter-satellite link load. These links become potential bottlenecks which can limit the overall backhaul performance even when the selected gateways have sufficient capacity. These results make routing a key challenge, as paths must be selected to avoid localized bottlenecks rather than simply minimize distance.

Preserving terrestrial services. Since the same LEO satellite network also carries terrestrial traffic, orbital backhaul should be carefully routed and isolated so that its traffic does not degrade existing terrestrial services. Efficient data delivery requires both adaptive gateway selection and load-aware routing, together with resource isolation between terrestrial and orbital traffic.

5 Conclusion

Satellites are emerging as a new class of network users whose growing backhaul demand cannot be met efficiently through direct ground station contacts alone. We envision reusing existing LEO satellite networks for delivering data from orbit to the ground. While promising, we examine that this vision requires rethinking access provisioning, link management, gateway selection, routing, and resource isolation. Looking forward, LEO networks may evolve from systems designed only for terrestrial users into a common communication substrate that also connects users in space.

Acknowledgment: We sincerely thank all anonymous reviewers for their valuable feedback. This work is supported by the National Natural Science Foundation of China (62572264), and the key fund of the National Natural Science Foundation of China (62132009).

References

[1] Blue Origin, LLC. 2026. Application for Authority to Launch and Operate the Blue Origin Orbital Data Center System. ICFS File No. SAT-LOA-20260310-00118. https://regmedia.co.uk/2026/03/20/fcc_filing.pdf Filed March 19, 2026.

[2] 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.

[3] Cloudflare Radar. 2026. SpaceX Starlink's Traffic Worldwide. https://radar.cloudflare.com/traffic/as14593. Accessed September 4, 2026.

[4] European Space Agency. 2017. Sentinel-2 Operations. https://www.esa.int/Enabling_Support/Operations/Sentinel-2_operations.

[5] European Space Agency. 2019. Working towards AI and Earth Observation. https://www.esa.int/Applications/Observing_the_Earth/Working_towards_AI_and_Earth_observation

[6] European Space Agency. 2020. Newcomers Earth Observation Guide. ESA Space Solutions. https://business.esa.int/newcomers-earth-observation-guide

[7] Ezra Feilden, Adi Oltean, and Philip Johnston. 2024. Why We Should Train AI in Space. Lumen Orbit (now Starcloud) White Paper, version 1.03. https://starcloudinc.github.io/wp.pdf

[8] International Telecommunication Union. 2011. *Handbook: Earth Exploration-Satellite Service*. ITU-R. https://www.itu.int/dms_pub/itu-r/opb/hdb/R-HDB-56-2011-PDF-E.pdf

[9] Gunter D. Krebs. 2026. Gunter's Space Page: Satellite Directory. https://space.skyrocket.de/directories/sat.htm.

[10] Yuanjie Li, Yimei Chen, Jiabo Yang, Jinyao Zhang, Bowen Sun, Lixin Liu, Hewu Li, Jianping Wu, Zeqi Lai, Qian Wu, and Jun Liu. 2025. Small-scale LEO satellite networking for global-scale demands. In *Proceedings of the ACM SIGCOMM 2025 Conference*. 917–931.

[11] Lixin Liu, Jinyao Zhang, Bijia You, Yimei Chen, Jiabo Yang, Yuanjie Li, Hewu Li, Qian Wu, Zeqi Lai, and Jun Liu. 2026. CrowdLink: Unlocking Idle LEO Network Capacity with User Terminals. In *1st New Ideas in*

Networked Systems (NINeS 2026). Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 28–1.

[12] NASA Jet Propulsion Laboratory. 2025. NISAR Press Kit. NASA Jet Propulsion Laboratory. https://www.jpl.nasa.gov/press-kits/nisar/

[13] Jason Rainbow. 2025. Starlink Mini Lasers to Link Muon Space Satellites for Near-Real-Time Connectivity. SpaceNews. https://spacenews.com/starlink-mini-lasers-to-link-muon-space-satellites-for-near-real-time-connectivity/

[14] Space Exploration Technologies Corp. 2026. Response to FCC Request for Clarification Regarding the SpaceX Orbital Data Center System. ICFS File No. SAT-LOA-20260108-00016. https://www.scribd.com/document/1045612919/SpaceX-Response-05-29-2026 Filed May 29, 2026.

[15] Space-Track.org. 2026. Space-Track.org. https://www.space-track.org/.

[16] SpaceX. 2022. Application for Fixed Satellite Service. FCC IBFS File No. SAT-AMD-20210818-00105. https://fcc.report/IBFS/SAT-AMD-20210818-00105

[17] SpaceX. 2026. Starlink Version 3 Satellites. https://starlink.com/updates/starlink-version-3-satellites

[18] Starcloud. 2026. Starcloud to Integrate SpaceX's Starlink Mini Lasers Into Its Orbital Data Center Constellation. Business Wire. https://www.businesswire.com/news/home/20260526670395/en/

[19] Starcloud, Inc. 2026. Application for Launch and Operating Authority for the Starcloud Orbital Datacenter System. ICFS File No. SAT-LOA-20260202-00073. https://www.scribd.com/document/993052648/Orbital-Data-Center-LOA-Narrative-1 Filed February 4, 2026.

[20] Bill Tao, Om Chabra, Ishani Janveja, Indranil Gupta, and Deepak Vasisht. 2024. Known knowns and unknowns: Near-realtime earth observation via query bifurcation in serval. In *21st USENIX Symposium on Networked Systems Design and Implementation (NSDI 24)*. 809–824.

[21] Deepak Vasisht, Jayanth Shenoy, and Ranveer Chandra. 2021. L2D2: Low latency distributed downlink for LEO satellites. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*. 151–164.

[22] Mike Wall. 2026. Another 'Star' Is Born: SpaceX Names AI Megaconstellation 'Starmind'. Space.com. https://www.space.com/space-exploration/launches-spacecraft/another-star-is-born-spacex-names-ai-megaconstellation-starmind

[23] Bingsen Wang, Xiaohui Zhang, Shuai Wang, Li Chen, Jinwei Zhao, Dan Li, and Yong Jiang. 2026. A large-scale ipv6-based measurement of the starlink network. In *IEEE INFOCOM 2026-IEEE Conference on Computer Communications*. IEEE, 1–10.

[24] Guanhua Wang, Fang Yang, Jian Song, and Zhu Han. 2024. Free space optical communication for inter-satellite link: Architecture, potentials and trends. *IEEE Communications Magazine* 62, 3 (2024), 110–116.

[25] Bo Wu, David Tipper, and Pengfei Zhou. 2025. Near-realtime earth observation via starlink leo satellite constellation. *arXiv preprint arXiv:2508.10338* (2025).