

CrowdLink: Unlocking Idle LEO Network Capacity with User Terminals

Lixin Liu ✉ 

Tsinghua University, Beijing, China

Bijia You ✉ 

Tsinghua University, Beijing, China

Jiabo Yang ✉ 

Tsinghua University, Beijing, China

Hewu Li ✉ 

Tsinghua University, Beijing, China

Zeqi Lai ✉ 

Tsinghua University, Beijing, China

Jinyao Zhang ✉ 

Tsinghua University, Beijing, China

Yimei Chen ✉ 

Tsinghua University, Beijing, China

Yuanjie Li¹ ✉ 

Tsinghua University, Beijing, China

Qian Wu ✉ 

Tsinghua University, Beijing, China

Jun Liu ✉ 

Tsinghua University, Beijing, China

Abstract

The Low Earth Orbit (LEO) network is booming worldwide thanks to its unprecedented number of satellites. However, most of these satellites remain underutilized to connect more users or boost performance, posing tensions for their return on investment. A critical cause is that their gateways to the Internet (ground stations) are geographically skewed or even centralized, forming last-mile bottlenecks. We examine the potential of eliminating these bottlenecks with ubiquitous user terminals (UTs). Our solution, CrowdLink, reuses UTs as *local access points* to decentralize satellites' gateways to the Internet, and as *relays* to convert idle satellite radio links into additional paths for more network capacity. This user-centric paradigm is self-scaling to more UTs and satellites (akin to P2P networks), resilient to rapid satellite mobility, mutually beneficial for users and operators, and readily deployable in operational LEO networks. Our real tests with Starlink UTs across three countries and large-scale simulations show that CrowdLink can increase each UT's throughput by $3.09\times$ on average (up to $65.27\times$), double the LEO network capacity utilization, and unlock 2.05–7.99 million more users for Starlink *without* adding satellites/ground stations.

2012 ACM Subject Classification Networks → Network architectures; Networks → Network design principles; Networks → Network protocol design

Keywords and phrases LEO Satellite Networks, User Terminal Relaying, Capacity Utilization

Digital Object Identifier 10.4230/OASICS.NINeS.2026.28

Funding This work is supported by the National Key Research and Development Plan of China (2022YFB3105201) and the National Natural Science Foundation of China (62572264).


Acknowledgements We thank our shepherd, Dr. Rachee Singh, and all reviewers for their valuable feedback.

1 Introduction

We are witnessing a boom in satellite Internet service. Low Earth Orbit (LEO) mega-constellations, such as Starlink, OneWeb, and Guowang, have deployed an unprecedented number of satellites to connect numerous users from rural, maritime, aviation, and other

¹ Corresponding author





■ **Figure 1** Reuse UTs as Internet gateways or relays to utilize idle satellites.


remote areas for Internet access everywhere, anytime. More LEO satellites are also planned to launch in the hope of expanding network capability for more customers and better performance.

However, this resource-intensive LEO network expansion has been increasingly questioned on its return on investment. Despite the noticeably decreased cost of manufacturing and launching each LEO satellite, the upfront cost for the entire mega-constellation with numerous LEO satellites remains prohibitive. Unless the LEO satellite mega-constellation can be efficiently utilized to serve enough customers, it would be tough for it to cover its cost for a breakeven (which is indeed the case for most LEO network operators today [31, 59, 58]).

Unfortunately, the current satellite network architecture hampers the full utilization of the LEO mega-constellation. As illustrated in Figure 1a, to provide Internet access, the LEO satellites need to connect to a terrestrial gateway (e.g., ground station) through either their local radio links or inter-satellite routes. In reality, the distribution of ground stations is more skewed and centralized than LEO satellites due to geographic/policy constraints, uneven demands for service, and prohibitive costs. On the one hand, these centralized ground stations become the last-mile bottleneck to throttle the overall LEO network capacity. On the other hand, most LEO satellites in motion are out of these ground stations' visibility, leaving their ultrahigh-capacity radio beams and laser crosslinks severely underutilized.

To eliminate this waste of satellites, an obvious solution is to decentralize the ground stations worldwide to avoid the last-mile bottlenecks. This method is impractical for at least two reasons. First, building such a global infrastructure is expensive. To fully utilize satellite links, each ground station requires dedicated phased array antennas, radio processing units, and fibers for broadband connectivity, leading to an upfront cost of \$1.25 million for just one Starlink community gateway [67]. The operators simply cannot afford to deploy a vast number of such ground stations for wide-area coverage. Second, ground stations cannot be deployed everywhere due to geographic constraints (e.g., in oceans that cover over 70% of the Earth) or policy reasons. This coverage blackhole can leave more than 70% LEO satellites still idle.

Instead, we study a complementary paradigm: *repurpose user terminals as satellites' Internet gateways and relays*. Compared to ground stations, satellite user terminals (UTs) are much cheaper (e.g., \$89–\$499 for each Starlink dish [1]) and more ubiquitously distributed worldwide at scale (e.g., over 7 million Starlink users from 150 countries, oceans, and aviation areas [64]). As shown in Figure 1b, they have built-in Ethernet ports to connect wired networks as local Internet gateways. Their antennas can also generate multiple beams to connect different satellites [2, 10, 50, 55, 28] to complement inter-satellite links for more capacity. While each UT's bandwidth is small, the aggregation of numerous UTs can result in a comparable bandwidth to dedicated ground stations. More importantly, this paradigm is *self-scaling* by nature (similar to P2P networks): more UTs can help themselves for more bandwidth and shorter path delays, while enabling the satellite operator to activate otherwise wasted satellite capacity and improve utilization at low cost.



■ **Figure 2** Underutilization of Starlink satellites based on Cloudflare [32].

Despite its appeal, this paradigm still confronts challenges from network scale and satellite mobility. Its numerous UT-based gateways/relays may inflate the network topology and complicate the satellite routing. This issue is exacerbated by the extreme mobility of LEO satellites. In addition, the fast-moving LEO satellites continually change their connectivity to different UT-based gateways/relays on a global scale. This implies frequent route reconfigurations that may downgrade network availability, efficacy, and reliability.

Our solution, CrowdLink, addresses these concerns with a simple UT-centric design. It leverages the fact that, compared to LEO satellites, most UTs can be viewed as almost stationary². The topological relations between UTs are hence stable as well despite satellite mobility. To this end, CrowdLink builds its routing and data forwarding over this stable UT mesh to mask the LEO dynamics for stable, scalable, and efficient networking. It instructs each UT to discover its upstream relay UT by geography, stabilizes each relay link by sharing the satellite schedules between neighboring UTs, constructs the data path to UT-based gateways by recursion, and incentivizes UTs to participate and forward data with built-in rewards. To ease incremental deployability, CrowdLink can realize all these features with readily-available modules in commodity UTs and operational LEO networks.

We have prototyped CrowdLink using Starlink UTs in a non-intrusive way and evaluated it with real-world experiments across three countries and large-scale simulations. Our results show that, without adding satellites or ground stations, CrowdLink can increase each UT's throughput by $3.09\times$ (up to $65.27\times$), double the LEO network capacity utilization, and unlock additional 2.05-7.99 million serviceable users for Starlink.


2 Motivation

This section measures the severe underutilization of LEO networks, analyzes the limitations of related work to address this issue, and motivates the use of UTs in our solution.

2.1 Idle Satellites in LEO Networks

LEO networks have expanded at an unprecedented pace in recent years. Due to each satellite's finite link capacity and coverage, most LEO networks tend to deploy more satellites for higher capacity, fewer coverage holes, and hence more customers and revenue. As of September

² The velocity of LEO satellites (about 27,000 km/h) is 2-4 orders of magnitude faster than airplanes (800-1,000 km/h), trains (up to 350 km/h), ships (20-50 km/h), vehicles (50-70 km/h), and human beings (4-7 km/h).



■ **Figure 3** The state-of-the-art LEO satellite network architecture.

2025, there have been over 11,856 LEO satellites in orbit [14], leaving fewer orbital slots to host more satellites. Such an orbital resource scarcity has further spurred the race of satellite deployments among LEO network operators, with up to 42,000 satellites to launch by Starlink [35, 36, 38], 3,232 by Amazon Kuiper [34, 37], and 13,000 by Guowang [15, 16].

Despite this crazy space race, most LEO satellites in orbit remain severely underutilized. Consider Starlink, the largest operational LEO network to date, with more than 7 million customers from 150 countries. With a 96-Gbps radio link to users and 3 laser inter-satellite links for each Starlink-v2 mini satellite [9], its overall network capacity has exceeded 450 Tbps [65]. However, according to Cloudflare’s traffic measurements for Starlink [32], more than 90% of its satellites’ bandwidth utilization is less than 20%, as illustrated in Figure 2b. Even in hotspots like New York, up to 26.6% of satellite capacity goes unused (Figure 2a). In low-demand areas (e.g., oceans), this number approximates 100%.


A main cause of this severe underutilization of satellites is their reliance on terrestrial gateways (ground stations) for Internet access. As illustrated in Figure 3, UTs’ traffic is first uplinked to satellites, optionally forwarded through inter-satellite links (ISLs), then downlinked to ground stations (GSs), and finally delivered to Internet via terrestrial points of presence (PoPs). However, large-scale GS deployment is capital-intensive and geographically constrained. For instance, Starlink currently operates only 227 ground stations and 46 PoPs globally [22], with most GSs centralized in few regions (Figure 4). In contrast, satellites are uniformly distributed over the globe. This spatial imbalance and scale asymmetry create a chokepoint: thousands of satellites must share access to a limited number of GSs. This results in the ground-to-satellite link (GSL) bottleneck, which limits per-UT throughput and reduces satellite utilization. Moreover, GSL bandwidth (10 Gbps [3, 4, 19]) is significantly lower than ISL bandwidth (200 Gbps [24, 9]), exacerbating GSLs as the dominant capacity bottleneck in the system.

2.2 Limitations of Related Work

To mitigate the above waste of satellites in LEO networks, the community has explored three categories of clean-slate designs or practical workarounds:

Expanding ground stations. A natural approach to alleviate GSL bottlenecks is to increase the number of ground stations and deploy them more widely. A denser GS deployment provides more landing opportunities for satellites, activating more GSLs, thereby increasing the utilization of idle satellite capacity. However, this approach encounters three challenges:

1. *Geographic and policy constraints:* GS deployment is primarily limited to land, which accounts for less than 30% of Earth’s surface. Even on land, many regions like rainforests, deserts, and mountainous areas can not easily support GS deployment [61]. Furthermore, geopolitical restrictions make it difficult to obtain the necessary permissions for GS installation, particularly in foreign or sensitive territories [57].




■ **Figure 4** Geo-distribution of Starlink's UTs, ground stations, and PoPs.

2. *High deployment cost:* Large-scale GS expansion incurs substantial economic overhead. The number of GSs must scale with the number of subscribers, leading to escalating costs, including site acquisition, construction, operational, and maintenance cost. For example, the Starlink community gateway incurs a recurring cost of \$75,000 per Gbps per month, with a one-time setup fee of \$1.25 million [67]. Similarly, deploying a single OneWeb ground station costs an estimated \$15 million, excluding operational and maintenance expenses [5]. As illustrated in Figure 6a, achieving just 11.23% satellite utilization improvement would require building an additional 795 GSs at a cost of 993.75–11,925 million USD, which would result in a low return on investment. Although it is possible to build cheaper ground stations [63], it is usually at the cost of limited bandwidth and hence falls short to eliminating the last-mile bottleneck.
3. *Inflexible:* Fixed GSs are inherently inflexible and cannot adapt to dynamic changes in user density and traffic demand. In practice, to meet peak demand, operators may oversupply GS capacity, which can reduce cost-efficiency during off-peak periods. [40, 13, 20, 17] suggest dynamically leasing GS capacity from third-party providers based on traffic demand, aiming to improve flexibility and reduce upfront deployment costs. However, this approach still faces two major limitations. First, the global GS pool remains small compared to the number of active satellites, and hundreds of thousands of satellites still compete for connections to only a few hundred GSs, leaving the fundamental bottleneck unresolved. Second, most GS providers [13, 20, 17] require satellite operators to reserve usage windows several hours in advance, which limits the system's responsiveness to short-term demand shifts.

Exploiting inter-satellite links. Alternatively, recent efforts seek to leverage ISLs to use idle satellites. These efforts typically involve designing inter-satellite topologies and load-balancing strategies to distribute user traffic more evenly across satellites and ground stations. As shown in Figure 5, ground station load is often imbalanced. Some GSs are heavily congested while others remain underutilized. By spreading traffic across multiple GSs, load balancing can improve per-UT bandwidth. Meanwhile, routing traffic through underutilized satellites helps increase satellite utilization.


While helpful, ISLs alone cannot eliminate last-mile GSL bottlenecks. When the dominant bottleneck is the aggregate GSL capacity, ISLs merely reshuffle where traffic lands and cannot increase the total throughput. Take Starlink as an example: with 227 GSs, each equipped with eight 20 Gbps Ka-band antennas, the system offers a total downlink capacity of 36.3




■ **Figure 5** Service capacity per GS.

■ **Table 1** Challenge for stable ISLs.

Laser terminal for inter-satellite link	Beam Divergence Angle
Mynaric CONDOR	15 μ rad[47]
TESAT SCOT80 (Telesat[12], Kepler[7])	>15 μ rad[11, 62]
GA-EMS Optical ISL LCT	30 μ rad[60]
CubeLCT	193 μ rad[56]



(a) Expanding ground stations.



(b) Exploiting inter-satellite links.

■ **Figure 6** Satellite utilization ratio of existing solutions.

Tbps. When shared by 4.6 million users, the average per-user capacity remains just 7.9 Mbps. Figure 6b illustrates the effect of ISL activation ratio on utilization. Assuming each satellite has no more than two intra-orbit ISLs and no more than one inter-orbit ISL, we simulate different ISL activation ratios. While utilization improves with more active ISLs, the overall gain is limited. Even at full ISL activation, the improvement is only 5.9%, which is far from sufficient.

In addition, ISLs remain operationally constrained in today's LEO networks. Establishing a laser ISL involves scanning, acquisition, tracking, and pointing, which can take hundreds of milliseconds [27, 66, 29]. Once connected, satellites must continuously adjust their orientation to maintain the connection. As shown in Table 1, ISL maintenance requires extremely tight angular alignment. Even small orbital jitter or line-of-sight disruptions can break the link. Hence, in operational LEO networks like Starlink, multi-hop ISL routing is mostly used for users outside the GS coverage [6].

Redesigning the LEO constellation. Since neither of the above workarounds can eliminate bottlenecks from skewed ground stations, it is interesting to consider whether one can totally redesign the LEO constellation to avoid satellite waste. In this direction, [45, 30] design non-uniform LEO constellations to match the satellite supplies with the uneven terrestrial demands. [53, 54, 48] advocate multi-operator LEO constellations to increase their satellite utilization. While inspiring for future LEO networks, these clean-slate constellation designs are difficult to retrofit into operational LEO networks today to utilize their massive idle satellites.

2.3 Opportunity: UTs as Ubiquitous Sidelinks

We unveil a new opportunity to unlock the LEO satellites' idle radio link capacity using ideas from P2P networks: *reuse UTs as ubiquitous side-links for satellites*. To date, modern satellite UTs are no longer dumb receivers. As summarized in Table 2, they are equipped

■ **Table 2** Built-in features of commodity satellite UTs that can be leveraged to enable UT-as-gateways/relays.

Representative UT	RF / Antenna	Local Connectivity	Software / Upgrade
Starlink Standard Kit	Electronic phased array; software assisted orienting[23]	Wi-Fi 6, tDSwo Ethernet LAN ports[23]	OTA software updates[18]
OneWeb Intellian OW11+CNX	Electronically scanned arrays; seamless connectivity[44]	Wi-Fi 6, multiple Gigabit ports[44]	OTA software updates[43]
Iridium MissionLINK® 700	High-gain electronic phased array; no moving parts[70]	Wi-Fi, Ethernet ports[70]	Firmware updates[69]

with phased array antennas for high-speed data links (which can form multiple beams for multi-satellite connectivity [50, 55, 28]), LAN interfaces to connect terrestrial networks, and over-the-air (OTA) firmware upgrades for various software-defined functions. These new features make UTs capable of playing two additional roles:

- **UT as a local access point (gateway):** Each UT can act as a lightweight, user-operated “mini ground station/PoP” by connecting to the terrestrial Internet via Ethernet, Wi-Fi, or other interfaces. Nearby UTs can reach the Internet through this UT gateway, providing an additional landing path when ground stations are congested or far away.
- **UT as a relay node:** If connected to multiple satellites using different radio beams, each UT can act as a relay between them. This side link can complement ISLs to increase capacity. It can also form hop-by-hop UT relay paths to extend the service coverage of the above UT-based local access points³.

By offloading traffic to these underutilized UT-Satellite links (USLs), this paradigm can utilize otherwise wasted satellite-to-ground capacity, increase each UT’s bandwidth, shorten its path delay, and improve the utilization of satellites to host more users. Compared to solutions in §2.2, it is more scalable and practical for four reasons:

1. **UTs are ubiquitous:** Unlike dedicated ground stations, UTs can be deployed wherever users reside or operate, including on oceans, islands, deserts, and mountainous regions. There have been more than 7 million Starlink UTs from 150 countries [64], with over 75,000 of them installed on ships worldwide [9]. This naturally forms a dense and flexible wide-area fabric of Internet access and relays, even in regions where ground stations are impractical to deploy. While each UT contributes limited bandwidth, the aggregation of massive UTs can proliferate bandwidth comparable to dedicated ground stations.
2. **UTs are cheap:** The above dense UTs are 3-4 orders of magnitude cheaper than dedicated ground stations and readily available, making them a cost-effective complement to state-of-the-art ground station networks.
3. **UTs are easy to upgrade:** Modern UTs already lay a solid foundation to incrementally deploy this paradigm, as shown in Table 2. For instance, Starlink Dishy already maintains a backup beam for inter-satellite soft handovers [2, 10, 28], which can be repurposed to maintain multi-satellite connectivities. In this way, the UT can act as an inter-satellite relay and complement ISLs for more bandwidth.
4. **Most UTs are stationary:** Compared to LEO satellites at about 27,000 km/h, the UT motions are negligible, making the overall UT relay/gateway fabric close to a fixed network. As we will see, this feature is especially helpful in scaling and simplifying the networking.

³ [41] discussed using UT relays as short-term replacements to ISLs for low-latency routing. But as pointed out by [42], ISLs still help with LEO network performance. This work is not intended to replace ISLs or ground stations. Instead, it complements them for more network capacity, higher satellite utilization (neither discussed in [41, 42]), and shorter delays.

3 Solution Overview

We propose CrowdLink, a user-centric scheme based on the vision in §2.3 to unlock the idle capacity in operational LEO networks. Although appealing, it is non-trivial to fulfill this vision due to various challenges for functionality, scalability, and sustainability. We discuss these challenges in §3.1 and present CrowdLink’s design principles to address them in §3.2.

3.1 Challenges

While UTs are already ubiquitously available and technically capable of leveraging satellites’ idle capacity, they still confront three challenges to fulfill their potential:

C1. Making it work: incremental deployment with minimal changes. We target a readily deployable design for today’s operational LEO networks. This constrains the solution space. Since satellites are harder to upgrade, the design should require no changes to satellites. Moreover, since current satellites assume ground stations as the only egress points, a key challenge is to integrate UT-based egress and relaying into the current framework while keeping the satellite side transparent. At the same time, given the scale of existing deployments, casting UTs as relays or local access points should not require hardware replacement.

C2. Making it work well: scalability under UT scale and LEO mobility. The scale of the LEO network will considerably inflate after incorporating numerous UT-based gateways and relays. Different from the state-of-the-art LEO networks, now the UTs are no longer end hosts only; they become part of the forwarding infrastructure and must be considered in routing and data forwarding at the scale of millions. More seriously, all these incorporated UTs’ links are intermittent due to LEO mobility. At the speed of 27,000 km/h, each LEO satellite can cover each UT for 3–10 minutes, after which the UT must hand over to another satellite to maintain connectivity. This results in tons of USL changes per second over this large-scale LEO network topology, thereby significantly challenging the scalability, efficiency, and reliability of routing and data forwarding.

C3. Making it work well continuously: incentives for UT participation. Unlike ground stations that are fully provisioned and managed by the network operators, UTs are privately owned and operated. While **C1** and **C2** highlight the feasibility and scalability challenges of enabling UTs as relay or breakout nodes, a further challenge lies in their willingness to participate. Providing such services is not free: a local breakout UT must contribute its own backhaul capacity and pay for Internet access and electricity, while a relay UT consumes additional device power and reduces its own usable capacity. Without incentives, participation may be insufficient. If incentives are misaligned or absent, the system risks either insufficient participation or opportunistic misuse. A practical design must incorporate an incentive mechanism that sustains continuous participation and adapts to spatial and temporal demand variations.

3.2 Key Ideas of CrowdLink

To address the challenges in §3.1, CrowdLink follows three design principles:




Figure 7 System overview of CrowdLink.

1. **Reuse existing network functions whenever possible:** CrowdLink minimizes deployment cost by reusing capabilities already present in operational LEO networks. In current systems, UTs are already managed by the operator for configuration and software updates, and the operator's controller routinely performs functions such as beam scheduling and gateway selection [68, 28, 52, 39]. CrowdLink builds on the existing control framework and extends its scope to incorporate UT-based local access points and relays. It does not require additional operator-side control primitives. On the UT side, CrowdLink relies on lightweight software upgrades to enable relay and local access point functions. As a result, existing UTs can be reused without hardware changes, and non-participating UTs remain unaffected. On the satellite side, it just treats the UT-based gateways as additional "ground stations" and UT relays as additional ISLs.
2. **UT-centric stable hierarchical relaying:** To mask LEO satellite dynamics, CrowdLink establishes a stable logical hierarchical topology among UTs. This leverages the fact that the topological relation between UTs is stable despite LEO satellite mobility. Each UT discovers its upstream UT (relay or local access point) based on geographic proximity, stabilizes this logical link by synchronizing its satellite schedules with this upstream UT, and recursively reaches the local access point UT along the resulting hierarchy. This avoids frequent routing updates that would otherwise be triggered by exhaustive USL changes and supports large-scale deployment via hierarchy.
3. **Geographic demand-aware UT rewards:** To secure long-term UT participations, CrowdLink incorporates a reward mechanism similar to P2P file sharing. Intuitively, a geographic area with higher bandwidth demands or lower satellite utilizations requires stronger incentives to attract local access points and relays. Hence, CrowdLink adopts a dynamic pricing scheme driven by these geographic supply-demand factors to reward UTs in this area and engage them for continuous contribution.

4 Design of CrowdLink

Figure 7 illustrates the workflow of CrowdLink. It is realized as an overlay on top of participating UTs and the operator's existing connectivity controller (which has already managed all UTs' satellite connectivity and handover schedules in Starlink [68, 28, 52, 39]). A user terminal can join CrowdLink by advertising its availability as a local gateway (if connected to the Internet) or relay node to the controller, or request the controller for side paths for more bandwidth or shorter delays. The controller maintains this pool of UT-based gateways and relays. Upon receiving a UT-side request for more paths, it allocates a set of UT-based gateways (with relays if needed) from this pool to fulfill this request, rewards them

with their satisfaction of such requests, and charges these requests. This generates additional revenues for both participating UTs and LEO network operators by utilizing their idle links, forming a sustainable ecosystem. We next elaborate on the technical details in each step.

4.1 Casting UTs as Internet Gateways/Relays

The first step for CrowdLink is to renovate commodity UTs as local access points (gateways) or relay nodes for satellites at minimal costs (**C1** in §3.1). To achieve so, it adopts a software-defined method for the UTs and connectivity controllers as follows:


UT-side upgrade. CrowdLink upgrades UT software to support two optional roles:

- *Local access point:* As shown in Table 2, commodity UTs are equipped with Ethernet ports or WiFi. They can use these interfaces to connect to the Internet and act as a local access point.
- *Inter-satellite relay:* [2, 10] show that many commodity UTs have backup beams to connect more than one satellite synchronously or asynchronously. CrowdLink leverages this capability, when present, to provide an additional satellite-to-satellite forwarding option through UTs. As shown in Figure 7, each relay UT can inspect incoming packets and tunnel them over the secondary beam. The UT acts as a transparent inter-satellite relay, without requiring any changes to satellite forwarding or routing logic.

Controller-side extension. Once the UTs can act as local access points and relay nodes, the operator-side connectivity controller should be capable of utilizing them. In real LEO networks today, the controller is already responsible for computing all UTs’ satellite schedules for its beams, i.e., the list of LEO satellites this UT’s primary beam should connect to over time (every 15s in Starlink [52, 51]). CrowdLink reuses this feature and applies to the secondary beam as well to generate its satellite schedules for relay UTs, ensuring that the second beam consistently connects to satellites required for relaying. This extension allows the controller to orchestrate dual-beam relays while preserving backward compatibility.

On-demand UT gateway/relay allocation. To request for a side path for more bandwidth and/or lower latencies, a UT can contact the controller to activate a secondary gateway. This procedure has been available in today’s LEO network, except that this secondary gateway referred to a dedicated ground station (PoP). CrowdLink reuses the same request workflow, but allows the controller to add a local access point UT as an additional egress option. For each request, the controller can select one or more UT-based gateways based on their location, bandwidth, and other factors. These UT gateways can co-exist with existing operator-managed PoPs: the requesting UT continue to rely on them as the baseline, while UT-operated gateways serve as supplementary performance boosters.

UT overlay-based transparent tunneling. For backward compatibility, CrowdLink implements the data tunneling across UTs without changing satellites. After selecting local access point UTs, the controller computes paths from normal UTs to its local access point UT via relay UTs (see details in §4.2). These paths are distributed to participating UTs. Each UT maintains a local path table keyed by (normal UT, local access point UT), which stores the ordered list of UTs along the path. Packets are then tunneled hop by hop according to these tables. Since all tunneling states are installed at the UT-side, satellites can be unaware of inter-UT relaying and forward traffic as usual.



■ **Figure 8** CrowdLink's transparent inter-UT traffic tunneling to satellites.


Next, we walk through CrowdLink's uplink and downlink data tunneling using this UT overlay:

- *Uplink tunneling:* Figure 8a shows how to tunnel traffic from the normal UT to its local access point. Packets originate at the normal UT, where CrowdLink intercepts outgoing traffic and decides whether to tunnel via a local access point UT. If so, the packet is encapsulated with a new routing header specifying the next-hop UT⁴. If multiple local access point UTs are assigned, packets can be distributed to leverage all paths. Relay UTs inspect the packet, update the header using their path tables, and tunnel it via secondary satellite to the next hop. Finally, the local access point UT verifies itself as the target, decapsulates the payload, and tunnels the original packet to the terrestrial Internet through its Ethernet port, maintaining NAT state for return flows.
- *Downlink tunneling:* Figure 8b shows the downlink process. Packets arriving from the terrestrial Internet are intercepted by the local access point UT, which consults its NAT table to identify the target normal UT. It then encapsulates the packet and tunnels it along the precomputed path distributed by the controller. Relay UTs update the routing header and tunnel accordingly. Upon receipt, the normal UT decapsulates the packet and delivers it to its local devices.

Secure and private tunneling. Since CrowdLink relies on user-operated, potentially untrusted UTs, it is understandable for users to be concerned about their security and privacy. The UT-based gateways and relays may eavesdrop on other UTs' data packets to leak privacy, manipulate them to violate integrity, or drop them to induce DoS attacks.

CrowdLink counteracts these threats with three measures. First, when a normal UT requests the controller for a side path, the controller will allocate an operator-controller PoP, instruct the allocated UT-based gateway to redirect traffic to this PoP, and establish an end-to-end secure tunnel between the normal UT and this PoP (which is already available

⁴ There are various options to implement this feature, such as IPv6's hop-by-hop header, SRv6, IP-in-IP encapsulation, or MPLS.



■ **Figure 9** Stable UT gateway-rooted relay tree via shared satellite schedule.

in Starlink [28] and hence reusable)⁵. This tunnel's end-to-end encryption and integrity protection can prevent intermediate relay/gateway UTs' eavesdropping and tampering. Second, to combat user behavior leakage from packet header metadata, CrowdLink can optionally use encrypted hop-by-hop packet encapsulation across multi-party UT relays and gateways, similar to the onion routing (Tor [33]). Third, CrowdLink penalizes misbehaving gateways and relays based on user feedback. After assigned a side path comprising UT gateways and relays, each normal UT routinely measures its performance (e.g., throughput, delay, and data loss) and reports these metrics to the controller. Based on these metrics from all these normal UTs, the controller can reliably detect underperforming or malicious UT relays/gateways, stop using them in future side-path assignments, and hence cut their rewards in §4.3 and block their misbehaviors.


Co-existence with legacy UTs, GS, and satellites. For legacy UTs, CrowdLink's gateway or next-hop relay UT is indistinguishable from another ground station/PoP. It can use them in the same way as using a secondary PoP. For ground stations, CrowdLink still keep them as the UTs' default gateway and complement them with relay/gateway UTs as side paths. For satellites, CrowdLink's UT gateways/relays are no different from a ground station. Its inter-UT tunneling is also transparent to satellites to facilitate their co-existence.

4.2 Stable Hierarchical inter-UT Relaying

While §4.1 enables UTs to act as relay nodes or local access points, the dynamic nature of LEO satellites (C2 in §3.1) introduces significant challenges. Links between functional UTs vary rapidly, forcing the controller to frequently reassign local access point UTs and recompute relaying paths. Such churn increases computation and signaling overhead, disrupts active sessions, and degrades service availability.

CrowdLink mitigates this issue by constructing a stable UT-level topology rooted at local access point UTs and progressively attaching relay UTs under schedule alignment constraints. Specifically, CrowdLink replicates schedule between parent and child UTs to align their satellite access whenever feasible. As a result, neighboring UTs can maintain access to a consistent satellite for long periods, despite satellite handovers caused by LEO mobility, thereby stabilizing the inter-UT logical link.

⁵ Note that these PoPs can still be decentralized across the Internet and not necessarily co-located with UT-based gateways to prevent bottlenecks.




■ **Figure 10** Aggregating UT gateway-rooted relay trees for multi-paths.

We next describe how CrowdLink builds this topology in practice: (i) constructing local access point trees that mask satellite dynamics; (ii) interconnecting these trees to enable multiple local access point connectivity; and (iii) inter-UT relaying over these trees.

Gateway-rooted inter-UT relay trees. CrowdLink organizes UTs into a forest of trees, each rooted at a local access point UT and composed of relay UTs and normal UTs as child nodes. For each newly added UT (*relay UT* or *normal UT*), the controller incrementally attaches it to an existing tree by performing the steps illustrated in Figure 9.

- *Step 1: Discover the upstream parent.* For a new UT, the controller searches for potential parents (local access point UTs or relay UTs already in a tree) within a radius of L , corresponding to the diameter of a satellite coverage area. For each candidate parent, the controller retrieves its satellite access schedule.
- *Step 2: Cross-check of satellite schedules.* For each candidate parent, the controller evaluates two metrics: (i) *schedule consistency*, defined as the duration over which the parent's access schedule remains available for the new UT (i.e., the time window during which both UTs can concurrently access the same satellite), and (ii) *load level*, measured as the number of existing downstream UTs already attached to the parent. Only candidates whose schedule consistency exceeds a threshold are retained. They are then ranked using a combined score that prioritizes local access point UTs over relay UTs, longer schedule alignment, and lower load levels. This heuristic balances stability and load distribution while avoiding unnecessary relay hops.
- *Step 3: Schedule replication.* The top-ranked candidate is selected as the parent. If the parent is a relay UT, its primary schedule is copied to the child (secondary beam for relay UT). If the parent is a local access point UT, its schedule is directly replicated to the child. This replication aligns the satellite access of the parent and child UTs, increasing the likelihood that their logical link remains available despite satellite mobility.

By repeating this process, the controller incrementally expands each tree outward from its local access point root. Once a normal UT joins a tree, it is implicitly associated with the local access point UT at the root. This hierarchical organization masks satellite dynamics, improves the utilization of underloaded UTs, and reduces the need for frequent path recomputation, thereby enhancing scalability and service continuity.



■ **Figure 11** Hierarchical relaying over CrowdLink's stable UT topology.

Aggregating inter-UT relay trees for multi-path. Because each local access point UT offers limited bandwidth and the distribution of downstream UTs may be uneven, CrowdLink further interconnects local access point-rooted trees to form a stable UT-level graph (Figure 10). This interconnection allows the controller to expose each normal UT to multiple reachable local access point UTs and to balance the load across them.


- *Step 1: Candidate discovery.* For each relay or normal UT, the controller searches within a radius L for local access point UTs or relay UTs in *other trees* and compares their schedules. Candidates whose schedule consistency exceeds a threshold are retained and ranked by the local access point UT of their tree, node type, and duration of alignment.
- *Step 2: Cross-tree interconnection.* The controller establishes cross-tree links among the selected candidates, prioritizing direct connections to local access point UTs from different trees to avoid unnecessary multi-hop relaying. Unlike intra-tree attachment, this process does not replicate schedules. Instead, the controller records these aligned links as additional forwarding options.

Through these steps, the controller constructs a stable UT-level graph in which local access point-rooted trees are interconnected by cross-tree links. This structure enables load-aware assignment of multiple local access point UTs to each normal UT, improving per-UT performance and increasing the utilization of otherwise wasted satellite capacity.

Hierarchical relaying across UTs. Given the stable UT-level graph, computing a path from a normal UT to a local access point UT reduces to a hierarchical tree traversal. Starting from the normal UT, the controller traverses upward along the intra-tree links until reaching the root local access point UT of the corresponding tree.

Figure 11 illustrates the process of relaying computation. When a normal UT (*Source UT*) initiates a service request, the controller attaches it to the tree rooted at local access point UT_1 and assigns UT_1 as its local access point. If the source UT requests additional bandwidth, the controller may allocate another local access point, such as UT_4 . Using aligned inter-tree links, the source UT is then treated as a leaf in the UT_4 -rooted tree. In each tree, the controller performs a *bottom-up traversal* from the source UT by repeatedly following stabilized parent links, yielding the corresponding relaying paths ($Source\ UT \rightarrow \dots \rightarrow UT_1$ and $Source\ UT \rightarrow \dots \rightarrow UT_4$).

Once the paths are determined, the controller distributes them to all involved UTs, including the source, intermediate relays, and local access points. Each UT installs the received path into its local path table for packet forwarding. Because UT-level links are stabilized through schedule alignment and replication (§4.2), these paths remain available for extended periods. As a result, CrowdLink avoids frequent path recomputation and reduces update signaling overhead, enabling scalable and highly available routing under LEO mobility.



■ **Figure 12** CrowdLink's geographic demand-aware incentive mechanism.

4.3 Geographic Demand-aware UT Rewards

To better activate idle LEO network capacity, CrowdLink is designed to encourage more UTs to participate and assist each other (**C3** in §3.1). Intuitively, geographic areas with higher bandwidth demands and/or lower satellite utilization require stronger incentives to attract nearby UTs to contribute. To this end, CrowdLink adopts a *geographic demand-aware rewarding mechanism* illustrated in Figure 12. Normal UTs pay for additional side paths, while relay and local access point UTs are rewarded proportional to their relayed/offloaded traffic based on dynamic prices. This mechanism builds on two core components: dynamic pricing and verifiable traffic accounting.


Dynamic pricing for traffic relaying. In CrowdLink, the connectivity controller assigns a per-unit reward price to each relay UT or local access point UT based on the local supply-demand condition. Both supply and demand are time-varying. Relay and gateway UTs may join or leave, and normal UTs' side-path requests can fluctuate (e.g., due to diurnal usage patterns). CrowdLink dynamically adapts prices to reflect these changes, so that relays and gateways in scarce or highly loaded areas tend to receive higher per-unit rewards, while rewards decrease when supply exceeds demand. This takes three steps:

- **UT-level load estimation.** The controller first estimates the load ratio ρ_i for each relay or local access point UT i : $\rho_i = \frac{load_i^{actual}}{load_i^{rec}}$, where $load_i^{rec} = \frac{b_i^{avail}}{b_{unit}}$ is the recommended load for i , b_i^{avail} is UT i 's contributed bandwidth, and b_{unit} is the estimated bandwidth required based on normal UTs' requests to the controller. $load_i^{actual}$ is the number of downstream normal UTs in the subtree rooted at UT i , capturing the amount of traffic that traverse it. $\rho_i > 1$ indicates UT i is overloaded and suggests that stronger incentive is needed.
- **Region-level load estimation.** To scale to millions of UTs, CrowdLink groups UTs' pricing by geographic regions. It partitions the Earth's surface into disjoint regions, each no larger than a satellite footprint. For each region R_k , the controller computes the average load ratio:

$$\bar{\rho}_k = \frac{1}{|R_k|} \sum_{i \in R_k} \rho_i$$

The base unit price for region R_k is then:

$$p_k^{base}(r) = p_0^{(r)} \cdot (1 + \alpha \cdot \max(0, \bar{\rho}_k - 1))$$



■ **Figure 13** CrowdLink’s proof-of-concept prototype over Starlink.

where $r \in \{\text{relay}, \text{local access point}\}$ specifies the UT’s role, and $p_0^{(r)}$ is the baseline cost for that role (with $p_0^{\text{local access point}} > p_0^{\text{relay}}$ to reflect higher costs such as Internet access). Using $\max(0, \bar{p}_k - 1)$ ensures that price inflation only occurs when regional demand exceeds sustainable capacity, avoiding over-incentivizing underutilized areas.

- *Contribution to topological connectivity.* Beyond traffic load, CrowdLink also adjusts rewards based on a UT’s contribution to connectivity. The goal is to reward UTs in positions that connect many normal UTs, especially in areas where relays or gateways are relatively scarce. Specifically, let d_i denote the path depth of UT i (the maximum hop count from a normal UT to i), and let s_i denote a normalized scarcity score derived from the fraction of normal UT paths traversing i and the number of equivalent UTs within the same tree. The final per-unit reward price for UT i is:

$$p_i = p_k^{\text{base}}(r_i) \cdot (1 + \beta \cdot d_i) \cdot (1 + \gamma \cdot s_i)$$

where β, γ are tunable coefficients. This formulation is intended to bias rewards toward UTs in high-demand regions and critical topological positions, so that incentives better reflect their contribution to connectivity. Based on this unit price, the connectivity controller pays each local access point UT (relay UT) for the traffic volume it offloads (relays).

Verifiable traffic accounting. An efficient incentive mechanism requires reliable and trustworthy measurement of forwarded traffic. CrowdLink deploys lightweight traffic accounting modules on all participating UTs. Each source, relay, and local access point UT locally records forwarded traffic and signs its report with verifiable credentials. At the end of each charging cycle, these signed records are submitted to the controller for consistency checks and settlement (e.g., using trust-free two-sided measurements [26], game-theoretical charging negotiation [46], or reputation-based charging [49]). The controller multiplies the settled traffic volume by the dynamic unit price p_i to compute (i) the bill charged to each normal UT and (ii) the rewards allocated to each relay and local access point UT. By combining geographic demand-aware pricing with verifiable accounting, CrowdLink aligns the incentives of normal UTs, serving UTs, and the LEO network operator, resulting in a triple win.

5 Proof-of-Concept Prototype

We showcase CrowdLink’s practicality with a proof-of-concept prototype over Starlink, the largest operational LEO network. As a third party, we cannot access Starlink’s network infrastructure or modify UT firmware. Therefore, our prototype adopts a *non-intrusive* design. We deploy a commodity host behind each Dishy and implement all CrowdLink functions on the host. This design requires no hardware modification and remains transparent to the satellite network. It allows us to validate CrowdLink’s packet relaying, controller orchestration, and runtime overhead in the wild.

Satellite-transparent UT relaying. For a source UT that requests a side path, the controller assigns a local access point UT and a chain of relay UTs to it. The host behind the source UT intercepts packets, applies controller-provided forwarding rules, and encapsulates packets. Each relay host decapsulates and reencapsulates packets before forwarding them further. The local access point host finally decapsulates packets and forwards them to the Internet through its terrestrial access link. Throughout this process, satellites operate as transparent forwarders and are unaware of inter-UT relaying. We achieve this design by allowing the controller to compute relay paths and enabling UT-side hosts to execute per-packet forwarding based on the computed results.


- **UT-side host:** Each host implements CrowdLink’s data-plane logic. It intercepts packets using Linux `netfilter/iptables` [25] and applies controller-assigned forwarding rules. For packets on relay paths, the host rewrites headers, performs encapsulation, and forwards them to the next-hop UT. To protect traffic from untrusted relay UTs, the source UT establishes an end-to-end encrypted tunnel with the operator-assigned egress endpoint. We use AES-128 as a showcase algorithm in evaluation to ensure payload confidentiality and integrity. In parallel with packet forwarding, the host records per-flow traffic volume on the `output` interface.
- **Controller:** The controller runs on a centralized ground server and implements the control-plane logic described in §4.2-§4.3. Upon receiving a service request, it selects local access point UTs and relay UTs and computes hierarchical relay paths over the stable UT-level topology. It then distributes relay paths to participating UTs and updates them when reconfiguration is required. At the end of each accounting cycle, it collects signed traffic records and generates bills and rewards.

Real-world deployment. As shown in Figure 13, we deploy three Starlink Dishys located in Portugal, Spain, and New York. They act as source UT, relay UT, and local access point UT, respectively. Each Dishy connects to a host running CrowdLink. This setup enables us to evaluate the functional correctness and performance of CrowdLink over real Starlink satellites.

6 Evaluation

We evaluate CrowdLink through a combination of proof-of-concept prototype and large-scale trace-driven simulation. Our main results are summarized as below:

- Our non-intrusive prototype in §5 can incrementally renovate commodity Starlink UTs as local access points and relays with unnoticeable performance overhead.
- By utilizing idle satellite links, CrowdLink increases each UT’s bandwidth by $3.09\times$ on average and up to $65.27\times$ compared to the state-of-the-art.
- By leveraging the state-of-the-art UTs, CrowdLink can double the overall capacity utilization of Starlink’s satellites.



■ **Figure 14** CrowdLink improves per-UT bandwidth and satellite capacity utilization.

6.1 Overall Effectiveness

This section uses large-scale simulations to evaluate whether UTs-as-relays/gateways can substantially improve the user experience and LEO network utilization.

Experimental setup. We evaluate CrowdLink using a trace-driven emulation based on Starlink, the largest operational LEO satellite network to date. Our emulation incorporates all active Starlink satellites and 227 ground stations. We model the user demand based on Starlink’s officially reported subscriber count (4.6 million) and the distribution of active users released in October 2024 [8]. We partition the Earth’s surface into disjoint $1^\circ \times 1^\circ$ latitude-longitude cells and assign users to cells accordingly. For each cell, we estimate the available satellite capacity based on satellite coverage, user distribution [8], and Starlink’s officially published regional per-UT data rates [21], assuming each satellite provides 96 Gbps radio link capacity [9]. The estimated unused GSL/USL capacity above each cell constrains the maximum bandwidth that relay UTs and local access point UTs within the cell can collectively contribute. To evaluate CrowdLink at scale, we proportionally sample 150,000 users according to the reported UT distribution. Among these sampled users, 50% are configured as local access points and 30% as relays. We compute per-UT bandwidth and satellite capacity utilization under varying activation ratios. We model the state-of-the-art LEO network architecture shown in Figure 3 as the *baseline*, where user traffic is uplinked to satellites and must be downlinked through ground stations before entering the terrestrial Internet.

UT-side bandwidth boost. Figure 14a shows the per-UT bandwidth improvement of CrowdLink over the baseline. In the baseline, all UTs compete for limited GSL capacity, so GSLs become the bottleneck and limit per-UT bandwidth, especially in hotspot areas. CrowdLink improves per-UT bandwidth by leveraging unused USL capacity to construct additional UT-based landing paths. A source UT can simultaneously use the default GS-based path and a controller-orchestrated side path that lands traffic through a local access point UT, thus increasing the total available bandwidth. When 10% of global UTs are active as sources, CrowdLink improves each UT’s bandwidth by $3.52\times$ (up to $21.88\times$). As more source UTs request side paths, multiple sources compete for the limited unused capacity that each local access point UT can contribute, so the average improvement gradually decreases. Even when all source UTs are active, CrowdLink still improves per-UT bandwidth by $3.22\times$ on average, with much larger gains ($65.27\times$) in hotspot areas where baseline throughput is severely constrained.

[Source UT ₁] [Portugal]	[Relay UT ₂] [Spain]	[Local access point UT ₃] [New York]
.098862 Capture outbound packet (dst = 8.8.8.8)	.120256 Capture packet (src = UT ₁ , dst = UT ₂)	.202755 Capture packet (src = UT ₂ , dst = UT ₃)
.098893 Look up path table (Next hop = UT ₂)	.120289 Look up path table (Next hop = UT ₃)	.202787 Look up path table (UT ₃ is the last hop)
.098910 Encapsulate packet (src = UT ₁ , dst = UT ₂)	.120555 Update packet header (src = UT ₂ , dst = UT ₃)	.202907 Decapsulate packet (dst = 8.8.8.8)
.098921 Transmit to next hop	.120655 Forward to next hop	.202970 Transmit to Internet

■ **Figure 15** Logs of CrowdLink’s proof-of-concept prototype.


Network-side capacity utilization. Figure 14b shows how CrowdLink improves the utilization of idle satellite capacity in Starlink. In the baseline, user traffic must be downlinked through ground stations. Due to the uneven distribution and limited capacity of satellite-ground links, a significant fraction of satellite capacity remains idle. We calculate the baseline utilization by measuring the aggregate bandwidth consumed by UTs in each region and comparing it to the aggregate satellite-ground link capacity, based on Starlink’s per-satellite satellite-ground capacity (96 Gbps [9]) and regional user distribution [8]. CrowdLink enables user traffic to land through local access point UTs through underutilized USLs, unlocking satellite capacity that would otherwise remain unused. When all UTs request side paths, CrowdLink nearly doubles the overall LEO network capacity utilization compared to the baseline. As more UTs participate in CrowdLink, a larger portion of idle satellite capacity is activated, leading to higher utilization.

6.2 Casting UTs as Internet Gateways/Relays


Next, we evaluate CrowdLink using our proof-of-concept prototype to validate the practical feasibility of reusing commodity UTs as relays and local access points in a real LEO network. Building on the non-intrusive, host-based prototype in §5, we deploy three Starlink Dishys located in Portugal, Spain, and New York, each connected to a host that implements relay or local access point functions. This setup allows us to examine whether UTs can correctly forward packets under real satellite dynamics and whether inter-UT relaying can sustain throughput comparable to native Starlink service.

Basic functionality. Figure 15 shows logs captured on the hosts behind each Starlink Dishy. The source UT encrypts the original payload, encapsulates it, and updates the destination address to the next relay UT. Upon receiving the packet, the relay UT rewrites the destination address to the controller-assigned local access point UT. Finally, the local access point UT decapsulates and decrypts the payload and forwards the original packet to the Internet. We adopt `tcpdump` tools at all UTs, confirming that the source UT successfully receives response packets from Internet servers, routed through the relay and gateway UTs. These results evaluate CrowdLink’s feasibility in practice.

UT performance boost. We compare the end-to-end throughput under UDP and TCP with/without CrowdLink. Figure 16 presents the results and reveals two key observations. First, UDP achieves relatively stable throughput, while TCP exhibits more fluctuations due to its sensitivity to packet loss and latency variation. Second, compared to baseline, CrowdLink significantly improves the data rate, reaching 24.91 Mbps in the uplink and 49.95 Mbps in the downlink. To understand the performance limits, we further examine the throughput bottleneck. As shown in Figure 16a, the uplink is the bottleneck in both baseline and CrowdLink, which limits the end-to-end throughput. This aligns with Starlink’s public measurements [21], which report an average global uplink rate of 24.95 Mbps. Since it



■ **Figure 16** CrowdLink improves end-to-end throughput.



■ **Figure 17** CrowdLink stabilizes relaying paths and reduces signaling overhead.


is largely independent of user or ground-station distribution, we attribute this to operator-imposed rate limiting. If such limits are relaxed in the future, CrowdLink can further increase throughput.

Packet processing overhead. We measure the packet processing latency introduced by CrowdLink to assess whether UT-side processing could become a performance bottleneck. Across all roles, the additional latency introduced by CrowdLink is on the order of tens to a few hundred microseconds. Specifically, the source UT adds $44.89 \mu\text{s}$ in the uplink and $40.60 \mu\text{s}$ in the downlink. A relay UT introduces an average processing latency of $246.74 \mu\text{s}$, while the local access point adds $189.29 \mu\text{s}$ in the uplink and $171.39 \mu\text{s}$ in the downlink. Compared to the end-to-end latency of LEO satellite networks, which is on the order of tens of milliseconds, this processing overhead is negligible and does not limit performance.

6.3 Stable Hierarchical inter-UT Relaying

We finally evaluate CrowdLink under large-scale simulation to examine whether its stable hierarchical inter-UT relaying can scale under extreme LEO mobility. We focus on three aspects that are critical in large-scale UT deployments: (i) relaying path stability and control overhead under frequent satellite handovers, (ii) scalability w.r.t. the number of relay UTs and local access point UTs, and (iii) the resulting end-to-end data performance.


Relaying path stability and signaling overhead. We first evaluate the stability of relaying paths and the associated signaling overhead. In the baseline, each relay UT independently schedules its secondary beam, so frequent satellite handovers directly translate into frequent




■ **Figure 18** Impact of the scale of relay UTs.

logical path changes. In contrast, CrowdLink organizes relay and local access point UTs as a tree rooted at local access point UT and stabilizes logical connectivity by replicating the parent UT's beam schedule to its children. Figure 17a shows the frequency of path updates from source UTs to their assigned local access point UTs under varying local access point ratios. Compared to the baseline, CrowdLink reduces the path update frequency by $2.63\text{--}3.12\times$. With more local access points, source UTs are more likely to establish direct connections, further reducing path dynamics. Since each path update requires the controller to push the updated routes to all UTs along the path, improved stability directly reduces signaling overhead. As shown in Figure 17b, CrowdLink reduces control messages by $5.9\text{--}9.0\times$.

Impact of the scale of relay nodes. We next examine how the number of relay UTs affects performance. Figure 18a shows the results when 10% of UTs act as local access points. Adding relay UTs increases the fraction of source UTs that can reach a local access point. However, since all traffic must ultimately exit through local access points, additional relays also intensify competition for limited egress capacity. As a result, the additional per-UT bandwidth improvement drops from $0.77\times$ to $0.64\times$, and the utilization of satellite capacity slightly drops from 24.91% to 24.03%. Figure 18b shows the case where 50% of UTs serve as local access points. Here, most source UTs already connect directly to local access points, thus relay UTs play a less critical role. Thus, increasing the number of relay UTs leads to only *mild* performance changes: the bandwidth improvement drops from $2.11\times$ to $2.06\times$, and utilization declines from 29.24% to 29.00%. These results indicate that relay UTs primarily extend connectivity, while network performance is ultimately constrained by the number of available local access points.



■ **Figure 19** Impact of the scale of local access point UTs.



■ **Figure 20** Transmission latency from source UTs to their gateways (GSs or local access point UTs).

Impact of the scale of local access points. We now evaluate how scaling the number of local access point UTs influences system performance. We tune the number of relay UTs to ensure that as many source UTs as possible can connect to at least one local access point, either directly or via relays. As shown in Figure 19, increasing the number of local access point UTs significantly improves both per-UT bandwidth and satellite capacity utilization. When all source UTs are active, the bandwidth improvement of CrowdLink over the baseline rises from $0.81\times$ to $4.21\times$, while the utilization of total contributed bandwidth increases from 23.47% to 31.65%. Unlike relay scaling, which shares a fixed set of exit UTs, each additional local access point UT contributes new uplink and downlink capacity and directly serves nearby users. This reduces contention and shortens paths, leading to higher per-UT throughput and better utilization of satellite resources.

End-to-end transmission latency. We finally compare the end-to-end latency from source UTs to their gateways. In the baseline, traffic is uplinked to satellites and downlinked through operator-managed ground stations. In CrowdLink, traffic is instead forwarded to user-controlled local access point UTs. Figure 20 shows the latency distribution. In our simulation, most source UTs offload traffic after a single satellite hop. This is because Starlink users are predominantly located near deployed ground stations, as reflected in the reported user distribution [8]. Since local access points are often geographically closer than fixed ground stations, CrowdLink achieves comparable to or lower latency than the baseline, reducing path delays by 6.8% (up to 22.97%). As the number of local access point UTs increases, more source UTs can establish direct connections, further shortening paths. These results indicate that CrowdLink preserves competitive end-to-end latency while providing more stable paths and higher capacity utilization.

7 Conclusion

This paper presents CrowdLink, a user-assisted solution to LEO network underutilization. CrowdLink aims to eliminate the last-mile bottlenecks and LEO satellite wastes induced by the skewed distribution of ground stations. Its key idea is to reuse user terminals as decentralized Internet gateways and relays for underutilized satellites, converting their idle links into complementary paths for more network capacity. This paradigm is self-scaling to massive users and satellites, mutually beneficial for them, and incrementally deployable in operational LEO networks. A key lesson from CrowdLink is that end users can play a valuable role in scaling, simplifying, and boosting the satellite network. We hope CrowdLink can inspire efforts for scalable and efficient Internet from space *for* the users, and *by* the users.

References

- 1 SpaceX Starlink. <https://www.starlink.com/>.
- 2 Starlink monitoring and hacking notes. <https://nelsonslog.wordpress.com/2021/03/10/starlink-monitoring-and-hacking-notes/>, 2021.
- 3 China's satellite achieves 10 Gbps space-to-ground laser data transfer. <https://interestinengineering.com/innovation/china-10-gbps-laser-data-transfer>, 2023.
- 4 Chinese satellite achieves high-Speed space-to-ground laser comms. <https://payloadspace.com/chinese-satellite-achieves-high-speed-space-to-ground-laser-comms/>, 2023.
- 5 Choosing OneWeb over Starlink: Can Taiwan build communications resiliency? <https://english.cw.com.tw/article/article.action?id=3865>, 2024.
- 6 Inter-satellite lasers are currently only used if the satellite cannot see the user terminal and ground station simultaneously. <https://x.com/elonmusk/status/1776985584596287622>, 2024.
- 7 Kepler validates Tesat optical communication terminals in orbit for Kepler network. <https://www.tesat.de/news/blog/942-kepler-validates-tesat-optical-communication-terminals-in-orbit-for-kepler-network>, 2024.
- 8 SpaceX starship update 2024. <https://www.youtube.com/watch?v=TUQzeUaGxBI>, 2024.
- 9 Starlink progress report 2024. https://www.starlink.com/public-files/starlinkProgressReport_2024.pdf, 2024.
- 10 Starlink REV4 now supports two beams? <https://x.com/olegkutkov/status/1785468069957005588?>, 2024.
- 11 Tesat SCOT80 optical communication terminals selected for GPS demonstration by Lockheed Martin. <https://www.tesat.de/news/press/1015-tesat-scot80-optical-communication-terminals-selected-for-gps-demonstration-by-lockheed-martin>, 2024.
- 12 Tesat selected by MDA Space as part of MDA Aurora™ supply chain to deliver optical inter-satellite links. <https://www.tesat.de/news/press/941-tesat-selected-by-mds-space-to-deliver-oi-sl>, 2024.
- 13 AWS ground station. <https://aws.amazon.com/ground-station/>, 2025.
- 14 CelesTrak. <https://celestrak.org/>, 2025.
- 15 China launches 8th batch of satellites for 13,000-strong internet megaconstellation (video). <https://www.space.com/space-exploration/launches-spacecraft/china-launches-8th-batch-satellites-guowang-satnet-internet-megaconstellation-video>, 2025.
- 16 e-submission of satellite network filings: CHN2025-73201. <https://www.itu.int/ITU-R/space/asreceived/Publication/DisplayPublication/60841>, 2025. Accessed on 2025-09-26.
- 17 Ground network services. <https://www.ksat.no/ground-network-services/>, 2025.
- 18 How do I update my Starlink if it doesn't automatically update? <https://www.starlink.com/gb/support/article/44b964f6-a538-d7c1-b893-b02822f444b5>, 2025. Accessed: 2025-09-29.

- 19 Singapore to test new intersatellite laser link tech in 2026. <https://payloadspace.com/singapore-to-test-new-intersatellite-laser-link-tech-in-2026/>, 2025.
- 20 Space connected, Limits gone. <https://leaf.space/>, 2025.
- 21 Starlink | availability map. <https://www.starlink.com/map>, 2025.
- 22 Starlink ground station locations (2025). <https://starlinkinsider.com/starlink-gateway-locations/>, 2025.
- 23 Starlink standard kit specification. <https://www.starlink.com/specifications>, 2025. Accessed: 2025-09-29.
- 24 Starlink technology. <https://www.starlink.com/us/technology>, 2025. Accessed on 2025-09-26.
- 25 The netfilter.org project. <https://www.netfilter.org/>, 2025.
- 26 SVR Anand, Serhat Arslan, Rajat Chopra, Sachin Katti, Milind Kumar Vaddiraju, Ranvir Rana, Peiyao Sheng, Himanshu Tyagi, and Pramod Viswanath. Trust-free Service Measurement and Payments for Decentralized Cellular Networks. In *Proceedings of the 21st ACM Workshop on Hot Topics in Networks (HotNets)*, pages 68–75, 2022.
- 27 Dhiraj Bhattacharjee, Aizaz U Chaudhry, Halim Yanikomeroglu, Peng Hu, and Guillaume Lamontagne. Laser inter-satellite link setup delay: Quantification, impact, and tolerable value. In *2023 IEEE Wireless Communications and Networking Conference (WCNC)*, pages 1–6. IEEE, 2023. doi:10.1109/WCNC55385.2023.10119016.
- 28 Chen Chen, Pavel Chikulaev, Sergii Ziuzin, David Sacks, Peter J Worters, Darshan Purohit, Yashodhan Dandekar, Vladimir Skuratovich, Andrei Pushkin, Phillip E Barber, et al. Low latency schedule-driven handovers, 2023. US Patent 11,729,684.
- 29 Weiqi Chen, Qi Zhang, Xiangjun Xin, Qinghua Tian, Ying Tao, Yufei Shen, Guixing Cao, Rui Ding, and Yifan Zhang. Beaconless acquisition tracking and pointing scheme of satellite optical communication in multi-layer satellite networks. In *Fifth Symposium on Novel Optoelectronic Detection Technology and Application*, volume 11023, pages 343–348. SPIE, 2019.
- 30 Yimei Chen, Lixin Liu, Yuanjie Li, Hewu Li, Qian Wu, Jun Liu, and Zeqi Lai. Unraveling physical space limits for LEO network scalability. In *Proceedings of the 23rd ACM Workshop on Hot Topics in Networks*, pages 43–51, 2024. doi:10.1145/3696348.3696885.
- 31 Tom Choi. OneWeb’s failure will dent new space investment. <https://interactive.sattelitoday.com/tom-choi-onewebs-failure-will-dent-new-space-investment/>, 2025.
- 32 Cloudflare Radar. SpaceX Starlink’s Traffic Worldwide. <https://radar.cloudflare.com/traffic/as14593>, year = Accessed: 2025-09-24.
- 33 Roger Dingledine, Nick Mathewson, and Paul Syverson. Tor: The second-generation onion router. In *13th USENIX Security Symposium*, 2004.
- 34 FCC. FCC authorizes Kuiper satellite constellation. https://docs.fcc.gov/public/attachments/FCC-20-102A1_Rcd.pdf, July 2020.
- 35 FCC. FCC grants SpaceX’s satellite broadband modification application. <https://docs.fcc.gov/public/attachments/FCC-21-48A1.pdf>, April 2021.
- 36 FCC. FCC partially grants SpaceX Gen2 broadband satellite application. <https://docs.fcc.gov/public/attachments/FCC-22-91A1.pdf>, June 2022.
- 37 FCC. Grant of Kuiper modification application to revise orbital parameters. <https://docs.fcc.gov/public/attachments/DA-24-224A1.pdf>, March 2024.
- 38 FCC. SpaceX authorized for SCS and operations at lower altitudes. <https://docs.fcc.gov/public/attachments/DA-24-1193A1.pdf>, December 2024.
- 39 Johan Garcia, Simon Sundberg, Giuseppe Caso, and Anna Brunstrom. Multi-timescale evaluation of starlink throughput. In *Proceedings of the 1st ACM Workshop on LEO Networking and Communication*, pages 31–36, 2023. doi:10.1145/3614204.3616108.
- 40 Chenwei Gu, Qian Wu, Zeqi Lai, Hewu Li, Yuxuan Weng, Weisen Liu, Jihao Li, Jun Liu, and Yuanjie Li. Novaplan: An efficient plan of renting ground stations for emerging LEO satellite networks. In *IEEE INFOCOM 2025 - IEEE Conference on Computer Communications*, pages 1–10, 2025. doi:10.1109/INFOCOM55648.2025.11044563.

- 41 Mark Handley. Using Ground Relays for Low-Latency Wide-Area Routing in Megaconstellations. In *Proceedings of the 18th ACM Workshop on Hot Topics in Networks (HotNets)*, pages 125–132, 2019. doi:10.1145/3365609.3365859.
- 42 Yannick Hauri, Debopam Bhattacharjee, Manuel Grossmann, and Ankit Singla. “Internet from Space” without Inter-satellite Links? In *Proceedings of the 19th ACM Workshop on Hot Topics in Networks (HotNets)*, pages 205–211, 2020. doi:10.1145/3422604.3425938.
- 43 Intellian. Flat panel compact series software release bulletin. https://downloads.ipinternational.net/oneweb/IntellianESA/RB-250611-01%20i-ESA_2.0.6.1ReleaseBulletin.pdf, 2025.
- 44 Intellian. OW11FL enterprise flat panel for land. https://go.intelliantech.com/l/350481/2025-09-25/2v4tsb/350481/1758830744aLFWY1ZZ/Intellian_OW11FL_Spec._Sheet_2025_DS0924_v1.9.pdf, 2025.
- 45 Yuanjie Li, Yimei Chen, Jiabo Yang, Jinyao Zhang, Bowen Sun, Lixin Liu, Hewu Li, Jianping Wu, Zeqi Lai, Qian Wu, et al. Small-scale LEO satellite networking for global-scale demands. In *Proceedings of the ACM SIGCOMM 2025 Conference*, pages 917–931, 2025.
- 46 Yuanjie Li, Kyu-Han Kim, Christina Vlachou, and Junqing Xie. Bridging the Data Charging Gap in the Cellular Edge. In *Proceedings of the ACM Special Interest Group on Data Communication (SIGCOMM)*, pages 15–28. ACM, 2019. doi:10.1145/3341302.3342074.
- 47 Jintao Liang, Aizaz U Chaudhry, Eylem Erdogan, Halim Yanikomeroglu, Gunes Karabulut Kurt, Peng Hu, Khaled Ahmed, and Stephane Martel. Free-space optical (FSO) satellite networks performance analysis: Transmission power, latency, and outage probability. *IEEE Open Journal of Vehicular Technology*, 5:244–261, 2023.
- 48 Lixin Liu, Yuanjie Li, Hewu Li, Jiabo Yang, Wei Liu, Jingyi Lan, Yufeng Wang, Jiarui Li, Jianping Wu, Qian Wu, Jun Liu, and Zeqi Lai. Democratizing Direct-to-Cell Low Earth Orbit Satellite Networks. In *21st USENIX Symposium on Networked Systems Design and Implementation (NSDI)*. USENIX, 2024.
- 49 Zhihong Luo, Silvery Fu, Mark Theis, Shaddi Hasan, Sylvia Ratnasamy, and Scott Shenker. Democratizing Cellular Access with CellBricks. In *ACM Special Interest Group on Data Communication (SIGCOMM)*, 2021.
- 50 Alireza Mehrnia, Masoud Kahrizi, Ahmad Mirzaei, Bagher Afshar, Ali SAJJADI, Bernd Pregardier, Ka Shun Carson PUN, NASIBY Omid, and Igor Elgorriaga. Channel extraction digital beamforming, 2025. US Patent US20240039568A1.
- 51 François Michel, Martino Trevisan, Danilo Giordano, and Olivier Bonaventure. A First Look at Starlink Performance. In *Proceedings of the 22nd ACM Internet Measurement Conference*, pages 130–136, 2022. doi:10.1145/3517745.3561416.
- 52 Nitinder Mohan, Andrew E Ferguson, Hendrik Cech, Rohan Bose, Prakita Rayyan Renatin, Mahesh K Marina, and Jörg Ott. A Multifaceted Look at Starlink Performance. In *Proceedings of the ACM Web Conference 2024*, pages 2723–2734, 2024. doi:10.1145/3589334.3645328.
- 53 Veronica Muriga, Swarun Kumar, Akshitha Sriraman, and Assane Gueye. A road map for the democratization of space-based communications. In *Proceedings of the 23rd ACM Workshop on Hot Topics in Networks*, pages 17–24, 2024. doi:10.1145/3696348.3696866.
- 54 Seoyul Oh and Deepak Vasishth. A call for decentralized satellite networks. In *Proceedings of the 23rd ACM Workshop on Hot Topics in Networks*, pages 25–33, 2024. doi:10.1145/3696348.3696896.
- 55 Eric Pepin, Kim Schulze, Amir Agah, Dean Carli, and Philippe Riondet. Systems and methods for beamforming in hybrid beamforming antennas, 2025. US Patent US12136977B2.
- 56 Patricia Martín Pimentel, Benjamin Rödiger, Christopher Schmidt, Christian Fuchs, Christoph Rochow, Thomas Hiemstra, Andreas Zager, Philipp Wertz, Marcus Knopp, Marc Lehmann, et al. Cube laser communication terminal (cubelct) state of the art. *Acta astronautica*, 211:326–332, 2023.

- 57 Tommi Rasila and Arto Ojala. *National Regulation of Satellite Ground Stations: A Global Comparison*, pages 195–217. Springer Nature Singapore, Singapore, 2024. doi:10.1007/978-981-97-3430-6_8.
- 58 Reuters. Starlink achieves cash-flow breakeven, says SpaceX CEO Musk. <https://www.reuters.com/technology/elon-musk-says-starlink-has-achieved-breakeven-cash-flow-2023-11-02/>, 2023.
- 59 Reuters. Ligado files for bankruptcy after stalled wireless expansion. <https://www.reuters.com/business/media-telecom/ligado-files-bankruptcy-after-stalled-wireless-expansion-2025-01-06/>, 2025.
- 60 David L. Robie, Aaron P. Freeman, and Robert E. Peterkin. MicroSat laser communication terminals and IR imaging payloads for space-based applications. In *Proceedings of the Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS)*. AMOS, 2020. URL: <https://amotech.com/TechnicalPapers/2020/Space-Based-Assets/Robie.pdf>.
- 61 Satsearch. Satellite ground stations: installing your own vs using an existing network. <https://blog.satsearch.co/2020-05-06-satellite-ground-stations-installing-your-own-vs-using-an-existing-network>, May 2020. Accessed: 2025-09-24.
- 62 SDA. Optical communications terminal standard version 4.0.0. https://www.sda.mil/wp-content/uploads/2024/07/SDA_OCT_Standard_4.0.0_final-20240701.pdf, 2024.
- 63 Vaibhav Singh, Akarsh Prabhakara, Diana Zhang, Osman Yağın, and Swarun Kumar. A Community-Driven Approach to Democratize Access to Satellite Ground Stations. In *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking (MobiCom)*, pages 1–14, 2021.
- 64 SpaceX. Starlink is connecting more than 7M people with reliable high-speed internet across 150 countries. <https://x.com/Starlink/status/1960867340951937239>, August 2025.
- 65 SpaceX. Starlink Network Update. <https://www.starlink.com/updates/network-update>, 2025.
- 66 Veronica Spirito, Ernesto Ciaramella, and Giulio Cossu. E2e physical layer and link analysis for high-throughput satellite optical communication. *IEEE Journal on Selected Areas in Communications*, 43(5):1660–1675, 2025. doi:10.1109/JSAC.2025.3545985.
- 67 Starlink. Community gateway. <https://www.starlinkinternet.info/community-gateway>, 2025. Accessed: 2025-09-24.
- 68 Hammas Bin Tanveer, Mike Puchol, Rachee Singh, Antonio Bianchi, and Rishab Nithyanand. Making Sense of Constellations: Methodologies for Understanding Starlink’s Scheduling Algorithms. In *Companion of the 19th International Conference on Emerging Networking EXperiments and Technologies (CoNEXT)*, pages 37–43, 2023. doi:10.1145/3624354.3630586.
- 69 Thales. Thales MissionLINK user guide. https://www.e-sat.fr/wp-content/uploads/2020/03/Thales_MissionLink_User_Manual.pdf, 2018.
- 70 Thales. Thales MissionLINK on Iridium Certus. <https://www.speedcast.com/wp-content/uploads/2021/11/Speedcast-Thales-Iridium-Certus-MissionLINK-C700-Product-Sheet.pdf>, 2021.