It’s a bird? It’s a plane? It’s CDN!

Investigating Content Delivery Networks in the LEO Satellite Networks Era

Rohan Bose
Technical University of Munich
Munich, Germany

Saeed Fadaei
University of Surrey
Guildford, UK

Nitinder Mohan
Delft University of Technology
Delft, Netherlands

Mohamed Kassem
University of Surrey
Guildford, UK

Nishanth Sastry
University of Surrey
Guildford, UK

Jörg Ott
Technical University of Munich
Munich, Germany

ABSTRACT

Content Delivery Networks (CDNs) have been pivotal in the dramatic evolution of the Internet, handling the majority of data traffic for billions of connected users. Low-Earth-Orbit (LEO) satellite networks, such as Starlink, aim to revolutionize global connectivity by providing high-speed, low-latency Internet to remote regions. However, LEO satellite networks (LSNs) face challenges integrating with traditional CDNs, which rely on geographical proximity for efficient content delivery – a method that clashes with the operational dynamics of LSNs. In this paper, we scrutinize the operation of CDNs in the context of LSNs, using Starlink as a case study. We develop a browser extension `NetMet` that performs extensive web browsing experiments from controlled nodes using both Starlink and terrestrial Internet access. Additionally, we analyse crowdsourced speed tests from Starlink users to Cloudflare CDN servers globally. Our results indicate significant performance issues for Starlink users, stemming from the misalignment between terrestrial and satellite infrastructures. We then investigate the potential for *SpaceCDNs* which integrate CDN infrastructure directly within the LSNs, and show that this approach offers a promising alternative that decreases latencies by over 50%, making them comparable with the CDN experience of users behind terrestrial ISPs. Our aim is to stimulate further research and discussion on overcoming the challenges of effective content delivery with growing LSN offerings.

CCS CONCEPTS

• **Networks → Network measurement**; **Network simulations**.

KEYWORDS

Starlink; LEO Satellite networks; CDN measurements; Internet Measurements

ACM Reference Format:
Rohan Bose, Saeed Fadaei, Nitinder Mohan, Mohamed Kassem, Nishanth Sastry, and Jörg Ott. 2024. It’s a bird? It’s a plane? It’s CDN!: Investigating Content Delivery Networks in the LEO Satellite Networks Era. In *The 23rd ACM Workshop on Hot Topics in Networks (HOTNETS ’24), November 18–19, 2024, Irvine, CA, USA.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3696348.3696879

1 INTRODUCTION

The Internet has evolved from a connection of a few endpoints to a “network of networks” interconnecting the world [11, 12, 24]. Simultaneously, the application traffic has become more content-centric fueled by the maturation of content delivery networks (CDNs) [37]. Currently, approximately 5.5 billion users contribute to generating nearly 330 exabytes of data daily, with CDNs handling nearly 70% of this traffic [7]. To cater to the the growing demand, CDNs have grown both *physically*, with expansive edge server deployments partnering with ISPs globally, and *logically* with sophisticated algorithms that map users to the “closest” server [6].

The recent emergence and growth of Low-Earth-Orbit (LEO) satellite mega-constellations, led by Starlink [46], OneWeb [38], etc., have introduced a new frontier in Internet connectivity – offering broadband-level connectivity to previously unreachable and underserved locations, such as remote rural areas and ocean-bound cruise ships. LEO satellite networks (LSNs) are witnessing a meteoric rise. For instance, the most extensive LSN, Starlink, now has over 3M+ subscribers in 100 countries covered by ≈ 6000 satellites [43], with plans to expand its fleet to 30,000 by 2027 [23, 42]. At these scales, LSNs are poised to become the first “global” ISPs. We argue that this places LSNs at odds with the common

---

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

ACM ISBN 979-8-4007-1272-2/24/11
https://doi.org/10.1145/3696348.3696879

CDN practice of geographically localising users and serving them from lowest latency servers [6, 19, 24]. Terrestrially-designed CDN technologies perform sub-optimally for users behind space-based ISPs, highlighting a fundamental divide between network optimization and geographic optimization.

CDNs inherently operate on the principle that geographic locality improves latency, using this to deliver geographically popular content to clients. However, in LSNs, the first point-of-contact with the terrestrial network is at the ground station (GS), which may/may not be geographically close to the client, especially when the connection traverses Inter Satellite Links (ISLs). As we discuss in §2, the ground station routes the connection terrestrially to its nearest LSN point-of-presence (PoP) which assigns a public IP address and connects the subscriber to the Internet. Consequently, content retrieval in LSNs relies on the LSN's terrestrial footprint, which undermines its ability to provide consistent connectivity globally [32, 36]. For instance, Starlink subscribers experience unwarranted geo-blocking from CDNs [17] when their connections are routed to PoPs deployed in countries where the requested content is geo-blocked [34].

In this work, we first investigate the state of the CDN performance for Starlink users globally through a suite of active measurements focusing on web browsing and passive data analysis of crowdsourced Cloudflare speed tests. Our analysis of over 1M+ measurements from 55 countries with Starlink coverage reveals significant degradation while retrieving region-specific content from CDNs due to ineffective terrestrial routing overheads. Based on these insights, we propose an integration of CDNs within the LSN infrastructure, a.k.a. *SpaceCDNs*, and explore potential consequences of such a network on content retrieval performance. Our attempt, with this work, is to motivate open discussion and further research to address substantial hurdles in delivering content effectively over LSNs, many of which we also highlight at the end of the paper. To foster reproducibility and enable future research, we publish our collected dataset and associated scripts at [4] and [3]

2 BACKGROUND AND RELATED WORK

Satellite Networking: A Primer. LSNs bear some similarity to traditional terrestrial network setups but also have some key differences. Starlink access is akin to cellular wireless, except the backhaul is satellite-based, and the wireless link operates on high-frequency Ku/Ka-band [25]. The most significant difference between terrestrial ISPs (including cellular operators) and LSNs, however, comes from the highly mobile *network infrastructure* of the latter. Satellites consistently revolve around the Earth at high-speeds reaching 27,000 km/h, almost 3× airplanes and 10× ground transportation [30, 32]. As of June 2024, Starlink operates 6,000+ satellites in altitudes ranging from 550 km to 1,200 km, with plans to launch 30,000+ satellites by 2027 in total including Very-Low Earth Orbits ($\approx$ 300 km) [46]. As a result, the connectivity between the user terminal (in Starlink ecosystem a.k.a. *Dishy*) and the satellite is constantly changing, with the satellite moving out of the line-of-sight within 5-10 minutes [26, 32].

To maintain consistent Internet access over such dynamically changing network topologies, LSN operators deploy ground stations (GS) that receive traffic from space back to the Earth, forming essentially a *bent-pipe* connection [40] (see fig. 1). Once the traffic reaches the ground, it is then routed to the traditional ground-based destination server through Points of Presence (PoP), that are strategically placed in datacenters and Internet eXchange Points (IXP) globally for direct access to the Internet backbone. Therefore, to operate a successful LSN, the operators must install extensive GSs worldwide in addition to launching a dense satellite constellation – mounting to a significant investment. To mitigate the need for numerous GSs, operators also use ISLs to route the traffic between satellites over long geographical distances to reach the ground station. For instance, Mohan et al. [36] found that Starlink users in the southern Africa connect to GSs and PoPs in Germany, nearly 9,000 km away, due to the lack of ground infrastructure in the region. Similar to a cellular core network, the internal traffic steering between terminal-to-GS is masked by a carrier-grade NAT and the client-side applications remain unaware of underlying network topology changes.

A Retrospective on CDN Operations. A content delivery network (CDN) is a hierarchy of geo-distributed servers designed to cache and serve content as close to the end-users as possible [37]. The primary goal of a CDN is to reduce bandwidth costs by minimizing WAN traffic and fulfilling user requests from the nearest cache server [5]. CDN operators achieve this objective through several strategies. Firstly, leading CDN providers, such as Akamai, Cloudflare, etc., partner with major ISPs globally to deploy servers within their networks (at the edge) or at the boundary of their network and Internet backbone (at PoP) [20]. Additionally, user requests are mapped to the "optimal" CDN cache based on network conditions and server load, using techniques like DNS-based redirection, anycast routing, and IP geolocation [6, 21].

Most internal CDN operations assume a static tree-like topology and user request influx from leaves of the hierarchy. For terrestrial ISP clients, this logical hierarchy maps closely to the physical geography – servers with (baseline) lower latencies are physically closer to users than those with higher latencies [16]. Satellite-based networks, however, completely disrupt this assumption, an example of which we illustrate in Figure 1. LSN client requests are always mapped

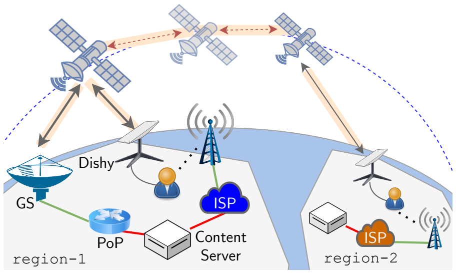

Figure 1: CDN reachability from terrestrial cellular ISP and LSN. Satellite subscribers connect to the CDN cache closes to the PoP, which results in a geographically distant server mapping for `region-2`.

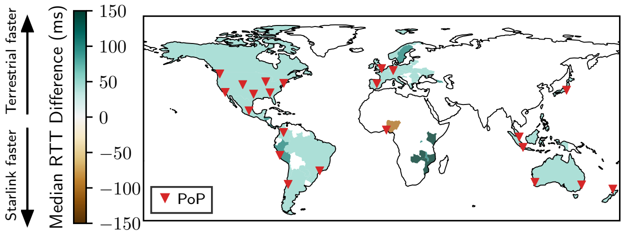

Figure 2: Delta difference (starlink - terrestrial) in Median RTTs to the most optimal CDN server location. Here we also illustrate the currently 22 operational Starlink PoP locations around the world.

to CDN servers nearest to the assigned PoP, which is comparable to terrestrial ISP if client resides in regions with local PoP deployment (e.g., `region-1`). However, for clients in `region-2`, the requests may be re-routed over series of ISLs to a geographically distant PoP, possibly located in a different country or even another continent [36]). Content fetching is suboptimal in such cases since the client skips the CDN server nearby and is mapped to a cache which likely lacks the `region-2` relevant content (e.g. news articles, videos, etc.). Many Starlink subscribers report challenges in accessing content from popular CDN-based services, such as Netflix, YouTube, etc., including geo-restrictions from other countries, slow loading times, and frequent buffering [17]. This disconnect in CDN mapping for LSNs is also costly for CDN operators as cache miss rates and content fetches over WANs are high for these users [10]. *Our key claim is that a* global ISP *such as Starlink that is not always able to find ground-based CDN infrastructure geographically close to the user (as §3 demonstrates) may benefit from placing caches directly on the satellites (§4).* This goes beyond proposals such as [14, 31] which use satellites as relays to terrestrial CDNs.

3 CONTENT DELIVERY OVER STARLINK

3.1 Measurement Methodology

To achieve an accurate representation of the global CDN performance over Starlink and the impact on application performance, we perform the following experiments.

Cloudflare AIM Dataset. We analyze the open-source Cloudflare Aggregated Internet Measurements (AIM) dataset [9], which includes speed test measurements from clients to Cloudflare CDN servers [8]. The speed test records typical metrics to assess user's Internet connection quality, such as download and upload speeds, latency, and jitter; but most importantly, provides an accurate representation of global Cloudflare CDN performance. We identify measurements from Starlink clients via their ASN (AS14593) and filter tests conducted since March 2024 – resulting in ≈ 22K measurements from 55 countries (almost 60% of Starlink coverage). For comparison, we also filter for measurements conducted by clients in terrestrial access networks. We carefully exclude (i) other satellite ISPs (e.g. Viasat, HughesNet) using Maxmind GeoIP DB [33] to ensure a fair comparison with Starlink, and (ii) networks that are not client-serving ISPs (e.g. cloud operators) using PeeringDB [41]. In total, we analyze ≈ 800K measurements from terrestrial clients in 196 countries. Note that we do not make a distinction between wired and wireless access as they are indistinguishable in the dataset. Since Cloudflare uses anycast to locate the "nearest" CDN server, we find that clients from the same city often target several CDN servers across different neighboring countries. We use the median of the *idle* latencies over both Starlink and terrestrial from a city to determine the "optimal" CDN server for the network at that location.

Web Browser Plugin. We develop a Chromium-based browser plugin, `NetMet` [1], that records real-time web browsing performance. `NetMet` periodically fetches the landing page of top-20 popular websites in the Tranco list [45] served by Cloudflare or Cloudfront CDN and records metrics such as DNS lookup time, TLS negotiation, server connect time, HTTP response time, etc. The plugin also identifies the CDN server location though `xhr-response` headers and the user location reported by the client's browser. We received measurements from volunteers using both Starlink and terrestrial ISPs in 8 and 15 countries, respectively, since March 2024 (totalling ≈ 5K measurements). Additionally, we dockerize the plugin and deploy it on LEOScope testbed probes [29] in UK, Germany, Canada and Nigeria to measure Starlink vs. terrestrial broadband from the same probe machine. Our containerized `NetMet` setup also records visual fidelity of web browsing, including metrics such as time-to-first-contentful-paint (FCP).

We make our collected passive and active measurement dataset and associated scripts available at [4] and [3].

3.2 Analysis

Global CDN performance. We first compare the idle latencies observed by Starlink users against those of terrestrial ISP users when connecting to the most optimal CDN location from a given city. Figure 2 shows the delta difference in median latencies (Starlink minus terrestrial) aggregated per country where measurements exist for both. Terrestrial connections almost always achieve lower latencies to CDNs, typically around 50 ms less than Starlink. The disparity is more pronounced for Starlink users in many African countries (e.g., Kenya, Mozambique, and Zambia), where latencies are around 120-150 ms higher. This is likely due to the large geographical distances (approximately 4000-8000 km) that must be covered, possibly through ISLs, as there are no nearby Starlink ground stations. As we discuss later, these connections are routed to the cache server via their assigned PoP in Frankfurt (Germany), also highlighted by other recent explorations [36, 40]. Table 1 highlights the average distances to the optimal CDN location and the corresponding median minimum latencies for selected countries over Starlink and terrestrial ISP connections. It is evident that in almost all regions, CDN access over Starlink is significantly worse than terrestrial alternatives, except for Spain, Japan and Germany which have locally deployed PoPs.

Figure 3 shows the case-study for users in Maputo (MPM), Mozambique in southern Africa accessing CDN servers over Starlink and terrestrial networks. The most frequent and optimal mapping over Starlink (see fig. 3a) is to the CDN in Frankfurt, which achieves the least observable latency ≈ 160 ms. When users connect to an African CDN (e.g. Cape Town), base latencies are higher due to the additional terrestrial path required after exiting the PoP towards the CDN, often exceeding 250 ms. Interestingly, we observe shorter latencies to other CDN locations in Europe (e.g. Lisbon) since the terrestrial network is well-provisioned in Europe. However, it is highly unlikely that the European CDNs will house content relevant for Africa, which will likely result in rerouting to origin servers over transcontinental WANs.

Table 1: The average geographical distance (in kms) to the best (= lowest latency) CDN server, indicating the sub-optimal CDN mapping for Starlink users.

Country	Terrestrial ISP		Starlink	
	Distance (km)	minRTT (ms)	Distance (km)	minRTT (ms)
Guatemala	6.9	7	1220.9	44.2
Mozambique	5.0	7.2	8776.5	138.7
Cyprus	34.7	7.45	2595.3	55.35
Swaziland	301.8	12.8	4731.6	122.7
Haiti	6.1	1.5	2063.2	50
Kenya	197.5	16	6310.8	110.9
Zambia	1202.64	44	7545.9	143.5
Rwanda	9.25	5	3762.8	87.5
Lithuania	168.6	12.4	1243.2	40
Spain	375.3	14.3	13.4	33
Japan	253	9	57.0	34

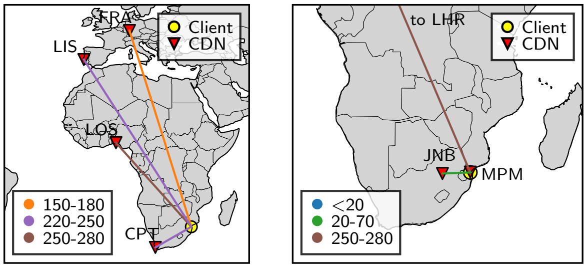

(a) Starlink ISP (b) Terrestrial ISP

Figure 3: Median latencies (in ms) to the connected Cloudflare CDN servers from Maputo, Mozambique.

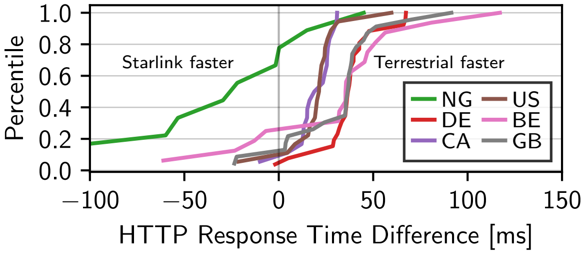

Figure 4: The difference in HTTP response times for Starlink and terrestrial ISPs from selected countries.

In contrast, Figure 3b illustrates a completely different CDN mapping for Maputo terrestrial ISP users. The most frequent connection goes to a CDN in Maputo itself, with latencies ≈ 20 ms; while other CDN locations in Africa (e.g. Johannesburg) are reachable within ≈ 70 ms. Note that for applications that care more about connecting to remote cloud servers, Starlink provides a faster and more reliable alternative with its fast-path to Europe; however, it falls rather short for fetching *geographically popular content.*

Web measurements. We now turn our attention to the web browsing performance from both Starlink and terrestrial collected by `NetMet` plugin. Figure 4 shows the HTTP response time (HRT) difference between the two networks. HRT is calculated as the time delay from when an HTTP request is made for a webpage to when the first byte of the response is received, excluding DNS lookup times and transport layer connection setup times, which might vary due to different end-user configurations. We find that the baseline latency differences discussed earlier also translate directly to the application-level performance. Terrestrial connections typically observe lower response times ≈ 20-50 ms, sometimes even up to 100 ms. Starlink users in Nigeria are the only

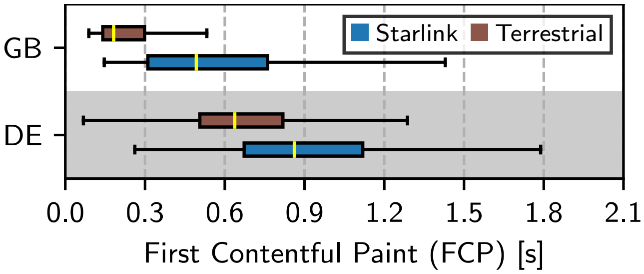

Figure 5: First Contentful Paint for Starlink and terrestrial ISPs in Germany (DE) and United Kingdom (GB).

outliers since they benefit from a nearby PoP and skip the still under-developed terrestrial infrastructure [18].

Figure 5 plots the first-contentful-paint (FCP) observed over both networks to provide deeper insight into user experience. FCP denotes time taken to render the first element on the page, and unlike response time, includes downloading necessary web elements from CDN servers. We only report experiments from Germany and United Kingdom since they represent the best-case scenario as both countries have local Starlink PoP deployment. Despite this, we find the median FCP values over Starlink are higher by ≈ 200 ms from both the locations. Considering the relative low and terrestrially-comparable baseline latencies over Starlink from these countries, user experience will be further degraded for countries situated geographically far from their assigned PoP (see Table 1). Latency-sensitive CDN-delivered web applications, such as live video streaming, video conferencing, cloud gaming etc., would suffer even further as Starlink suffers from significant *bufferbloat* [36]; which we also corroborate as we observed > 200 ms during active downloads from (potentially) ISL-enabled countries, e.g. Costa Rica, Bulgaria, Kenya, etc. Our results hint at several possible benefits if the satellites were to also cache content objects in space, which would reduce the access latencies and dependence on PoPs. We explore such an integration next.

4 CDN & LSN: A MATCH MADE IN SPACE

In this section, we explore the "obvious" *what-if* scenario of locating CDN caches on LEO satellites that are directly overhead of users. Clearly this would lead to greatly decreased latencies. The biggest challenge to overcome is the dynamics of LSNs — the fact that LEO satellites can move out of range of a given user location within minutes (*cf.* §2). (We leave to §5 discussion of other challenges such as managing the power budget; making the economics "work", etc.).

Figure 6 illustrates one possible solution to the problem of maintaining continuous CDN cache service to a user despite LEO dynamics: An object can be fetched from the satellite directly overhead, providing it is in the cache. If not, ISLs can be used to fetch the object from the nearest satellite which caches it. If no close-by satellite has the object cached, the request is sent to a cache server on the ground.

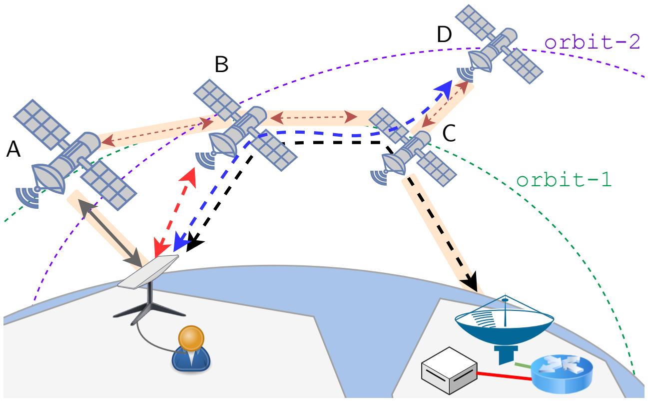

Figure 6: SpaceCDN overview. (i) Content is cached on satellites, allowing the client to fetch it directly (red arrow). (ii) If the content is not on the nearest satellite, ISL links route the request to the next closest satellite with the cached content (blue arrow). (iii) If content is unavailable on all satellites, the request is routed back to the cache server near ground station (black arrow).

Longer requests, such as streaming video, can take advantage of the predictable movements of LSNs. For example, a video object can be striped (correlating to a collection of DASH segments) such that the first stripe of n minutes is cached on the first satellite if it will be visible to the user for the first n minutes of playback; the next few stripes can be located on the second satellite which will be overhead of the user while its stripes are being served, and so on. Note that while Stripe 1 is being streamed to the user by a satellite A, subsequent stripes can be uploaded onto the caches of the satellites such as B and C that follow, thereby hiding the latency of the bent-pipe. Satellites in LSN orbits revisit a location roughly every 90 minutes; so cached video stripes can be accessed by other users if they are directly under the caching satellite, or using ISLs to fetch the nearest copy.

The success of such an approach depends on the relative latency improvements that can be achieved. We estimate this using `xeoverse`, a recently released simulator, which is calibrated with Starlink for realistic latencies and other channel characteristics under different conditions [27]. We configure `xeoverse` to simulate Shell 1 of Starlink (1,584 Satellites). Figure 7 shows the CDF of the latency to fetch objects from a satellite cache $n = 1, 2, 3, 5, 10$ ISL hops away, comparing it with the latencies from Starlink and terrestrial ISPs via AIM measurements (Table 1 shows the lowest observed latency; here we plot the whole CDF). It is evident that even on terrestrial ISPs, CDN access latency has a long tail.

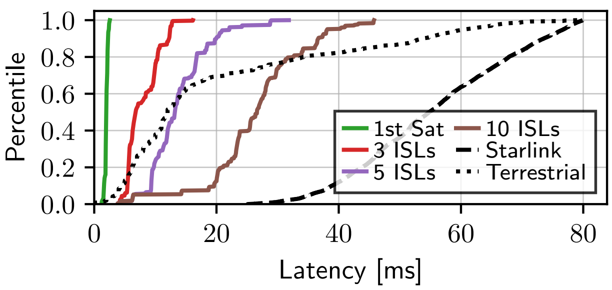

Figure 7: Starlink-CDN (dashed black) and Terrestrial ISP-CDN (dotted black) latencies obtained from AIM data (§3.1) compared with SpaceCDN Latencies (solid colours) obtained from `xeoverse` [27] simulation. The different colors show the latencies when content is found within 1, 3, 5 and 10 ISL hops. SpaceCDN is competitive with terrestrial ISPs if content is found within 5 hops, but when comparing with current Starlink latencies, even 10 ISL hops offers around half the latency.

If objects can be fetched in five ISL hops or fewer, LSNs can offer comparable performance to CDNs connected to terrestrial ISPs, even outperforming them in the tail of the distribution. This is mainly because of the low latencies achievable over the fully optical ISLs. < 5 hops is eminently doable in modern LSNs. As a case in point, Starlink Shell 1 consists of 72 orbital planes, with 22 satellites in each plane. Thus, with around 4 copies distributed within each plane, an object can be reachable within 5 hops, even within a single orbital plane; fewer copies would be needed if east-west ISLs across orbital planes are also used.

5 DISCUSSION

While a conjoined operation of CDNs and LEO satellite networks (LSNs) shows promise, it also opens several research questions and challenges, which we discuss in this section.

Operational overheads of cache servers in satellites. Satellites are engineered to be lightweight and power-efficient to keep launch costs low and operational lifetime high. Therefore, it is important to understand the impact of installing cache servers on satellite thermal control, power management and operational performance. For answering this, we turn towards recent works exploring the feasibility of in-orbit edge computing [2, 13]. The authors in [2] find that supporting a high-end server inside a Starlink satellite is not prohibitive in terms of weight and volume. The same high end-server comes with an attached storage capacity of ≈ 150 TB [15]. Considering, Starlink currently operates 6000 satellites (and more in future), the total storage capacity that the satellite constellation might be able to host, will be upwards of 900 PB i.e. > 300M 2-hour long 1080p videos at 30FPS. Xing et al. [47] study the performance of commercial off-the-shelf hardware within Cubesat satellites (≈ 16% size and weight of Starlink satellites) by launching it into space. The authors find that while the additional power consumption is manageable by cyclical solar energy generation, continued computation might degrade the on-board battery lifespan. Moreover, since the satellites are passively cooled and must remain below 30°C to maintain safe operations, the heat generated during active computation/content caching also raises concerns of (potential) thermal issues. However, the authors find that the overall temperature only exceeds the threshold after hours of continuous computation, which can be mitigated by intelligent request scheduling. Since caching is not computationally intensive, SpaceCDNs are likely more feasible than in-orbital computing, but will need managing.

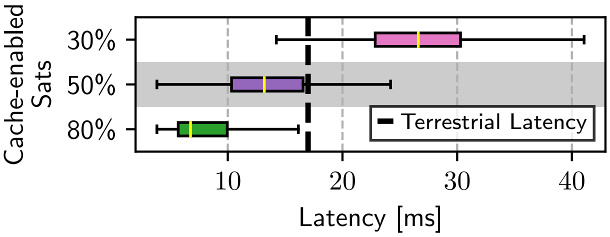

Figure 8: SpaceCDN latencies when only 30%, 50% and 80% of satellites are duty-cycling as caches, with the rest serving as relays to reach the caches. The vertical line indicates the median latency achieved from terrestrial ISPs to content CDN; showing SpaceCDNs are competitive with ≥ 50% satellites caching at a time.

A potential first-cut approach might be to duty-cycle the satellites caches: each satellite would only periodically serve content; at other times, it would relay its users' content requests over ISLs to nearby satellites which are currently offering a caching service. To check the feasibility of this, we rerun our previous simulation, randomly picking x% of the whole fleet of satellites as caches in each duty cycle slot.

Figure 8 shows the latencies observed when a fraction x = 80%, 50% and 30% of satellites are serving as content caches at any given point of time. The horizontal line indicates the median latency from terrestrial ISPs to CDNs. It can be seen that even when only half (50%) of satellites are currently duty-cycling as caches and the rest are relays, SpaceCDNs can be competitive with terrestrial ISP-CDN latencies. However, thermal and power management challenges require further in-depth investigation before SpaceCDNs become practical.

Content Bubbles Given the predictable nature of both the satellite orbits and content popularity in different geographical regions (e.g., a Boca Juniors vs River Plate soccer game is likely to be popular mostly over South America, or more specifically, in Argentina), we foresee the potential of machine learning algorithms to predict and prefetch content on satellites as they approach field-of-view of a country. We envision new algorithms will be needed to form such localized *content bubbles* forming over different geographical regions where the infrastructure moves but the content remains accessible. For instance, a satellite moving from over the US to Europe can use content-aware cache eviction to eliminate American Football and pre-fetch soccer content.

Space VMs Today's terrestrial CDNs go far beyond caching and enable stateful applications to be served from the edge. CDNs today are critical for low-latency use cases, such as coordinating state across users within a local area in multiplayer games [35]. Enabling such applications can be trickier with SpaceCDNs because the satellites are in constant motion. In future work, we plan to explore the possibility of locating replicated VMs on successive satellites that will be serving a geographic area, and use techniques developed for VM migration [28, 48] in data centers to sync the state change deltas ($\approx < 100$ MBs) from the satellite currently serving an area to the satellite(s) which will be overhead next, thereby providing seamless operations to the users within their coverage area.

Economics of Space CDNs. Table 1 shows that benefits of SpaceCDNs is most pronounced in regions with poor terrestrial Internet connectivity, for instance parts of Africa and South America. However, these regions are also not the most lucrative markets for network operators, which raises the question of how SpaceCDNs play into the traditional CDN operations. We envision a MetaCDN-like model [22] where the LSNs own and operate their satellite caches (possibly partnering with existing *local* terrestrial CDN operators) and allow multiple customers (e.g. streaming services) to cache their content on the satellites. The benefits of utilizing SpaceCDNs go far-beyond catering to LSN-connected users but can even improve content reachability. For instance, content providers can leverage the natural trajectory of satellite caches to distribute geographically-relevant content without the traversing either WAN or ISL links – opening dimensions for *content wormholing*. Similarly, recent direct-to-cell cooperations between LSNs and terrestrial ISPs necessitate several other avenues for SpaceCDN-catered content delivery to maintain seamless user experience [44].

Expansion of LSN ground infrastructure. Starlink currently operates approximately 150 ground station and 20 PoP locations all around the world [39], and continues to expand its coverage to more countries. However, the rate of expansion of its ground infrastructure is still much slower than its satellite launches. The primary cause of this is bureaucratic roadblocks in the form of spectrum licensing, land acquisition, setting up and consequent maintenance of this infrastructure (including PoPs). Additionally, the terrain and remote conditions of certain regions globally may not be suitable for the construction and maintenance of ground infrastructure – notwithstanding, the several thousand km of cabling required to connect Starlink's ground network to the internet backbone. Keeping aside all these issues, even with sufficient and steady ground infrastructure expansion, we only foresee the best case latency to hover around 20-30ms (as observed in §3). In this regard, SpaceCDNs appear as a viable opportunity as they not only provide significant latency reduction (which may match or even outperform terrestrial alternatives) but also reduce the continuous maintenance overhead of the operators.

6 CONCLUSION

The rapid growth of Low-Earth Orbit (LEO) satellite networks (LSNs), such as Starlink, has the potential to revolutionize global Internet connectivity. However, they are at odds with traditional Content Delivery Networks (CDNs) since they rely on ground stations for connectivity, leading to non-optimal server mapping and geographical restrictions for users in remote regions connecting via inter-satellite links. Of course, Starlink will make Internet access better for those in remote areas. However, the aim should not be to "settle" for this improved access, but rather, see if there are feasible solutions to put them on an equal footing with those in well connected urban areas, with access to high speed terrestrial broadband and established CDN server presence.

In this paper, we first showcased the disparity in CDN performance for Starlink users globally through a comprehensive suite of active measurements and passive data analysis, and identified several challenges in content delivery over LSNs. We then explored the potential for *SpaceCDNs* by integrating CDN infrastructure directly within the LEO satellites and ground stations via extensive simulations. We demonstrated that SpaceCDNs offer a promising extension to traditional CDN operations and identified several open research directions for effective content delivery over LSNs.

ACKNOWLEDGEMENT

This research was supported by the National Growth Fund through the Dutch 6G flagship project "Future Network Services", the German Federal Ministry of Education and Research joint project 6G-life (16KISK002) and the German Federal Ministry for Digital and Transport project 5G-COMPASS (19OI22017A). We thank Cloudflare for making their measurements dataset publicly available and several volunteer users of the `NetMet` plugin who graciously ran measurements from their networks on our behalf.

REFERENCES

[1] 2024. *NetMet: Internet Performance Check Tool.* https://github.com/boserohan/netmet

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

[3] Rohan Bose, Saeed Fadaei, Nitinder Mohan, Mohamed Kassem, Nishanth Sastry, and Jörg Ott. 2024. *It’s a bird? It’s a plane? It’s CDN!: Investigating Content Delivery Networks in the LEO Satellite Networks Era: Code.* https://github.com/boserohan/starlink-cdn-analysis

[4] Rohan Bose, Saeed Fadaei, Nitinder Mohan, Mohamed Kassem, Nishanth Sastry, and Jörg Ott. 2024. *It’s a bird? It’s a plane? It’s CDN!: Investigating Content Delivery Networks in the LEO Satellite Networks Era: Dataset.* https://doi.org/10.14459/2024mp1756495

[5] Matt Calder, Ryan Gao, Manuel Schröder, Ryan Stewart, Jitendra Padhye, Ratul Mahajan, Ganesh Ananthanarayanan, and Ethan Katz-Bassett. 2018. Odin: Microsoft’s Scalable Fault-Tolerant CDN Measurement System. In *15th USENIX Symposium on Networked Systems Design and Implementation (NSDI 18)*. USENIX Association, Renton, WA, 501–517. https://www.usenix.org/conference/nsdi18/presentation/calder

[6] Fangfei Chen, Ramesh K. Sitaraman, and Marcelo Torres. 2015. End-User Mapping: Next Generation Request Routing for Content Delivery. In *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication* (London, United Kingdom) *(SIGCOMM ’15)*. Association for Computing Machinery, New York, NY, USA, 167–181. https://doi.org/10.1145/2785956.2787500

[7] Inc. Cisco Systems. [n. d.]. Cisco Annual Internet Report (2018–2023) White Paper. https://www.cisco.com/c/en/us/solutions/collateral/executive-perspectives/annual-internet-report/white-paper-c11-741490.html. Accessed: 2024-06-26.

[8] Cloudflare. [n. d.]. Internet Speed Test - Measure Network Performance. https://speed.cloudflare.com/. Accessed: 2024-06-26.

[9] Cloudflare. 2023. Cloudflare: Measuring network quality to better understand the end-user experience. https://blog.cloudflare.com/aim-database-for-internet-quality/.

[10] Cloudflare. 2023. Cloudflare Radar - Starlink traffic. https://blog.cloudflare.com/radar-2023-year-in-review#globalstarlink. Accessed: 2024-06-26.

[11] Lorenzo Corneo, Maximilian Eder, Nitinder Mohan, Aleksandr Zavodovski, Suzan Bayhan, Walter Wong, Per Gunningberg, Jussi Kangasharju, and Jörg Ott. 2021. Surrounded by the Clouds: A Comprehensive Cloud Reachability Study. In *Proceedings of the Web Conference 2021* (Ljubljana, Slovenia) *(WWW ’21)*. Association for Computing Machinery, New York, NY, USA, 295–304. https://doi.org/10.1145/3442381.3449854

[12] The Khang Dang, Nitinder Mohan, Lorenzo Corneo, Aleksandr Zavodovski, Jörg Ott, and Jussi Kangasharju. 2021. Cloudy with a chance of short RTTs: analyzing cloud connectivity in the internet. In *Proceedings of the 21st ACM Internet Measurement Conference* (Virtual Event) *(IMC ’21)*. Association for Computing Machinery, New York, NY, USA, 62–79. https://doi.org/10.1145/3487552.3487854

[13] Bradley Denby and Brandon Lucia. 2020. Orbital Edge Computing: Nanosatellite Constellations as a New Class of Computer System. In *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems* (Lausanne, Switzerland) *(ASPLOS ’20)*. Association for Computing Machinery, New York, NY, USA, 939–954. https://doi.org/10.1145/3373376.3378473

[14] Alexis A. Dowhuszko, Juan Fraire, Musbah Shaat, and Ana Pérez-Neira. 2020. LEO satellite constellations to offload optical terrestrial networks in placement of popular content in 5G edge nodes. In *2020 22nd International Conference on Transparent Optical Networks (ICTON)*. 1–6. https://doi.org/10.1109/ICTON51198.2020.9203447

[15] Hewlett Packard Enterprise. 2024. HPE ProLiant DL325 Gen10 Server. https://www.hpe.com/psnow/doc/a00045498enw?jumpid=in_lit-psnow-red. Accessed: 2024-10-01.

[16] Xun Fan, Ethan Katz-Bassett, and John Heidemann. 2015. Assessing affinity between users and CDN sites. In *Traffic Monitoring and Analysis: 7th International Workshop, TMA 2015, Barcelona, Spain, April 21-24, 2015. Proceedings 7*. Springer, 95–110.

[17] Nicole Feist. 2022. I Tried Elon Musk’s Starlink Internet on a Royal Caribbean Cruise Ship. https://www.royalcaribbeanblog.com/2022/06/28/i-tried-elon-musks-starlink-internet-royal-caribbean-cruise-ship. Accessed: 2024-06-26.

[18] Agustin Formoso, Josiah Chavula, Amreesh Phokeer, Arjuna Sathiaseelan, and Gareth Tyson. 2018. Deep Diving into Africa’s Inter-Country Latencies. In *IEEE INFOCOM 2018 - IEEE Conference on Computer Communications*. 2231–2239. https://doi.org/10.1109/INFOCOM.2018.8486024

[19] Benjamin Frank, Ingmar Poese, Yin Lin, Georgios Smaragdakis, Anja Feldmann, Bruce Maggs, Jannis Rake, Steve Uhlig, and Rick Weber. 2013. Pushing CDN-ISP collaboration to the limit. *SIGCOMM Comput. Commun. Rev.* 43, 3 (jul 2013), 34–44. https://doi.org/10.1145/2500098.2500103

[20] Benjamin Frank, Ingmar Poese, Yin Lin, Georgios Smaragdakis, Anja Feldmann, Bruce Maggs, Jannis Rake, Steve Uhlig, and Rick Weber. 2013. Pushing CDN-ISP collaboration to the limit. *SIGCOMM Comput. Commun. Rev.* 43, 3 (jul 2013), 34–44. https://doi.org/10.1145/2500098.2500103

[21] Palak Goenka, Kyriakos Zarifis, Arpit Gupta, and Matt Calder. 2022. Towards client-side active measurements without application control. *SIGCOMM Comput. Commun. Rev.* 52, 1 (mar 2022), 20–27. https://doi.org/10.1145/3523230.3523234

[22] Oliver Hohlfeld, Jan Rüth, Konrad Wolsing, and Torsten Zimmermann. 2018. Characterizing a meta-CDN. In *Passive and Active Measurement: 19th International Conference, PAM 2018, Berlin, Germany, March 26–27, 2018, Proceedings 19*. Springer, 114–128.

[23] Bin Hu, Xumiao Zhang, Qixin Zhang, Nitin Varyani, Z. Morley Mao, Feng Qian, and Zhi-Li Zhang. 2023. LEO Satellite vs. Cellular Networks: Exploring the Potential for Synergistic Integration. In *Companion of the 19th International Conference on Emerging Networking EXperiments and Technologies* (, Paris, France,) *(CoNEXT 2023)*. Association for Computing Machinery, New York, NY, USA, 45–51. https://doi.org/10.1145/3624354.3630588

[24] Geoff Huston. 2023. The Future Internet Through the Lens of History. https://www.comsoc.org/publications/ctn/future-internet-through-lens-history. Accessed: 2024-06-26.

[25] Jayasuryan V Iyer, Khasim Shaheed Shaik Mahammad, Yashodhan Dandekar, Ramakrishna Akella, Chen Chen, Phillip E Barber, and Peter J Worters. 2022. System and method of providing a medium access control scheduler. US Patent 11,540,301.

[26] Liz Izhikevich, Manda Tran, Katherine Izhikevich, Gautam Akiwate, and Zakir Durumeric. 2024. Democratizing LEO Satellite Network Measurement. , 26 pages. https://doi.org/10.1145/3639039

[27] Mohamed M Kassem and Nishanth Sastry. 2024. xeoverse: A Real-time Simulation Platform for Large LEO Satellite Mega-Constellations. *IFIP Networking*.

[28] Horacio Andrés Lagar-Cavilla, Joseph Andrew Whitney, Adin Matthew Scannell, Philip Patchin, Stephen M Rumble, Eyal De Lara, Michael Brudno, and Mahadev Satyanarayanan. 2009. Snowflock: rapid virtual

machine cloning for cloud computing. In *Proceedings of the 4th ACM European conference on Computer systems*. 1–12.

[29] LEOScope. [n. d.]. A Measurement Testbed for Low-Earth Orbit (LEO) Satellite Networks. https://leoscope.surrey.ac.uk/. Accessed: 2024-06-26.

[30] Yuanjie Li, Hewu Li, Wei Liu, Lixin Liu, Wei Zhao, Yimei Chen, Jianping Wu, Qian Wu, Jun Liu, Zeqi Lai, and Han Qiu. 2023. *A Networking Perspective on Starlink's Self-Driving LEO Mega-Constellation*. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3570361.3592519

[31] Michele Luglio, Simon Pietro Romano, Cesare Roseti, and Francesco Zampognaro. 2019. Service Delivery Models for Converged Satellite-Terrestrial 5G Network Deployment: A Satellite-Assisted CDN Use-Case. *IEEE Network* 33, 1 (2019), 142–150. https://doi.org/10.1109/MNET.2018.1800020

[32] Sami Ma, Yi Ching Chou, Haoyuan Zhao, Long Chen, Xiaoqiang Ma, and Jiangchuan Liu. 2022-12-28. Network Characteristics of LEO Satellite Constellations: A Starlink-Based Measurement from End Users. (2022-12-28). arXiv:2212.13697 [cs.NI]

[33] MaxMind. [n. d.]. GeoIP Databases. https://www.maxmind.com/en/geoip-databases. Accessed: 2024-06-26.

[34] Allison McDonald, Matthew Bernhard, Luke Valenta, Benjamin VanderSloot, Will Scott, Nick Sullivan, J. Alex Halderman, and Roya Ensafi. 2018. 403 Forbidden: A Global View of CDN Geoblocking. In *Proceedings of the Internet Measurement Conference 2018* (Boston, MA, USA) *(IMC '18)*. Association for Computing Machinery, New York, NY, USA, 218–230. https://doi.org/10.1145/3278532.3278552

[35] Chou Mo, Guowei Zhu, Zhi Wang, and Wenwu Zhu. 2018. Understanding Gaming Experience in Mobile Multiplayer Online Battle Arena Games. In *Proceedings of the 28th ACM SIGMM Workshop on Network and Operating Systems Support for Digital Audio and Video* (Amsterdam, Netherlands) *(NOSSDAV '18)*. Association for Computing Machinery, New York, NY, USA, 25–30. https://doi.org/10.1145/3210445.3210450

[36] Nitinder Mohan, Andrew E. Ferguson, Hendrik Cech, Rohan Bose, Prakita Rayyan Renatin, Mahesh K. Marina, and Jörg Ott. 2024. A Multifaceted Look at Starlink Performance. In *Proceedings of the ACM on Web Conference 2024* (Singapore, Singapore) *(WWW '24)*. Association for Computing Machinery, New York, NY, USA, 2723–2734. https://doi.org/10.1145/3589334.3645328

[37] Erik Nygren, Ramesh K. Sitaraman, and Jennifer Sun. 2010. The Akamai network: a platform for high-performance internet applications. *SIGOPS Oper. Syst. Rev.* 44, 3 (aug 2010), 2–19. https://doi.org/10.1145/1842733.1842736

[38] OneWeb. 2023. OneWeb: Connect with Ease. https://oneweb.net/connect_with_ease. Accessed: 2023-05-26.

[39] Nathan Owens. 2024. Unofficial Crowdsourced Starlink Global Gateways PoPs. https://bit.ly/3tykCGW. Accessed: 2024-06-26.

[40] Jianping Pan, Jinwei Zhao, and Lin Cai. 2024. Measuring the Satellite Links of a LEO Network. In *IEEE International Conference on Communications*.

[41] PeeringDB. [n. d.]. PeeringDB. https://www.peeringdb.com/. Accessed: 2024-06-26.

[42] Aravindh Raman, Matteo Varvello, Hyunseok Chang, Nishanth Sastry, and Yasir Zaki. 2023. Dissecting the Performance of Satellite Network Operators. *Proc. ACM Netw.* 1, CoNEXT3, Article 15 (nov 2023), 25 pages. https://doi.org/10.1145/3629137

[43] Starlink. 2023. Starlink coverage map. https://www.starlink.com/map.

[44] Starlink. 2023. *Starlink Direct to Cell.* https://direct.starlink.com/ Accessed: 2023-10-12.

[45] Tranco. [n. d.]. A Research-Oriented Top Sites Ranking Hardened Against Manipulation. https://tranco-list.eu/. Accessed: 2024-06-26.

[46] Wikipedia. 2023. Starlink. https://en.wikipedia.org/wiki/Starlink.

[47] Ruolin Xing, Mengwei Xu, Ao Zhou, Qing Li, Yiran Zhang, Feng Qian, and Shangguang Wang. 2024. Deciphering the Enigma of Satellite Computing with COTS Devices: Measurement and Analysis. In *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking* (Washington D.C., DC, USA) *(ACM MobiCom '24)*. Association for Computing Machinery, New York, NY, USA, 420–435. https://doi.org/10.1145/3636534.3649371

[48] Fei Zhang, Guangming Liu, Xiaoming Fu, and Ramin Yahyapour. 2018. A survey on virtual machine migration: Challenges, techniques, and open issues. *IEEE Communications Surveys & Tutorials* 20, 2 (2018), 1206–1243.