



A Multifaceted Look at Starlink Performance

Nitinder Mohan*
Technical University of Munich
Germany

Andrew E. Ferguson*
The University of Edinburgh
United Kingdom

Hendrik Cech*
Technical University of Munich
Germany

Rohan Bose
Technical University of Munich
Germany

Prakita Rayyan Renatin
Technical University of Munich
Germany

Mahesh K. Marina
The University of Edinburgh
United Kingdom

Jörg Ott
Technical University of Munich
Germany

ABSTRACT

In recent years, Low-Earth Orbit (LEO) mega-constellations have ushered in a new era for ubiquitous Internet access. The Starlink network from SpaceX stands out as the only commercial LEO network with over 2M+ customers and more than 4000 operational satellites. In this paper, we conduct a first-of-its-kind extensive multi-faceted analysis of Starlink performance leveraging several measurement sources. First, based on 19.2M crowdsourced M-Lab speed tests from 34 countries since 2021, we analyze Starlink global performance relative to terrestrial cellular networks. Second, we examine Starlink's ability to support real-time latency and bandwidth-critical applications by analyzing the performance of (i) Zoom conferencing, and (ii) Luna cloud gaming, comparing it to 5G and fiber. Third, we perform measurements from Starlink-enabled RIPE Atlas probes to shed light on the last-mile access and other factors affecting its performance. Finally, we conduct controlled experiments from Starlink dishes in two countries and analyze the impact of globally synchronized "15-second reconfiguration intervals" of the satellite links that cause substantial latency and throughput variations. Our unique analysis paints the most comprehensive picture of Starlink's global and last-mile performance to date.

CCS CONCEPTS

• **Networks** → **Network measurement.**

KEYWORDS

Starlink; Satellite communications; Internet measurements

ACM Reference Format:

Nitinder Mohan, Andrew E. Ferguson, Hendrik Cech, Rohan Bose, Prakita Rayyan Renatin, Mahesh K. Marina, and Jörg Ott. 2024. A Multifaceted Look at Starlink Performance. In *Proceedings of the ACM Web Conference 2024 (WWW '24)*, May 13–17, 2024, Singapore, Singapore. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3589334.3645328>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WWW '24, May 13–17, 2024, Singapore, Singapore

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0171-9/24/05...\$15.00

<https://doi.org/10.1145/3589334.3645328>

1 INTRODUCTION

Over the past two decades, the Internet's reach has grown rapidly, driven by innovations and investments in wireless access [22, 46, 47] (both cellular and WiFi) and fiber backhaul deployment that has interconnected the globe [3, 8, 10, 24, 78]. Yet, the emergence of Low-Earth Orbit (LEO) satellite networking, spearheaded by ventures like Starlink [66], OneWeb [50], and Kuiper [4], is poised to revolutionize global connectivity. LEO networks consist of mega-constellations with thousands of satellites orbiting at 300–2000 km altitudes, promising ubiquitous *low latency* coverage worldwide. Consequently, these networks are morphing into "global ISPs" capable of challenging existing Internet monopolies [67], bridging connectivity gaps in remote regions [36, 70], and providing support in disaster-struck regions with impaired terrestrial infrastructure [21].

Starlink from SpaceX stands out with its expansive fleet of 4000 satellites catering to 2M+ subscribers across 63 countries [60, 77]. The LEO operator plans to further amplify its coverage and quality of service (QoS) by launching $\approx 42,000$ additional satellites in the coming years [15]. However, despite significant global interest and the potential to impact the existing Internet ecosystem, only limited explorations have been made within the research community to understand Starlink's performance. The challenge stems from a lack of global vantage points required to accurately gauge the network's performance since factors such as orbital coverage, density of ground infrastructure, etc., can impact connectivity across regions. Initial studies have resorted to measurements from a handful of geographical locations [25, 35, 36, 40] or extrapolated global performance through simulations [26] and emulations [31]. However, the community agrees on the limited scope of such studies and has made open calls to establish a global LEO measurement testbed to address this challenge [52, 61, 74]. Some researchers have navigated around this hurdle by exploring alternative measurement methods, e.g., by targeting exposed services behind user terminals [19], by mining measurements on social media platforms [73], or by recruiting users in select regions [58]. While innovative, we argue that these techniques are insufficient to uncover the intricacies affecting the network, specifically its capability to support applications.

This paper addresses this knowledge gap and provides the first comprehensive multi-faceted measurement study on Starlink. Our work is distinct from previous research in several ways. Firstly,

* Authors contributed equally to this research.

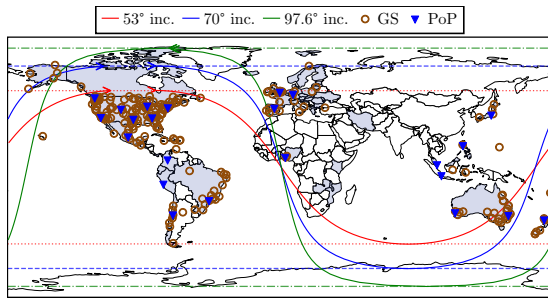


Figure 1: Orbits of three Starlink inclinations and crowd-sourced Ground Station (GS) and Point-of-Presence (PoP) locations [51]. Shaded regions depict Starlink’s service area.

we examine the global evolution of the network since 2021 by analyzing the M-Lab speed test measurements [12] from 34 countries (largest so far). We complement our investigation through active measurements over 98 RIPE Atlas [59] probes in 21 countries and conduct high-resolution experiments over controlled terminals in two European countries to investigate real-time web application performance and factors impacting Starlink’s last-mile access. Specifically, we make the following contributions in this work.

(1) We present a longitudinal study of global Starlink latency and throughput performance from M-Lab users in §4. Our analysis, incorporating ≈ 19.2 M samples, reveals that Starlink performs competitively to terrestrial cellular networks. However, its performance varies globally due to infrastructure deployment differences, and is dependent on the density and closeness of ground stations and Points-of-Presence (PoPs). We also observe signs of *bufferbloating* as Starlink’s latency significantly increases under traffic load.

(2) We assess and compare the performance of real-time web applications, specifically Zoom video conferencing and Amazon Luna cloud gaming, to terrestrial networks (§5). We find that, under optimal conditions, Starlink is capable of supporting such applications, matching the performance over cellular; however, we do observe some artifacts due to the network’s periodic reconfigurations.

(3) We perform targeted measurements from Starlink RIPE Atlas [59] probes and leverage their diverse locations to characterize the satellite last-mile performance (§6.1). We find that the “bent-pipe” (terminal \leftrightarrow satellite \leftrightarrow ground station) latency within the dense 53° shell remains consistent worldwide (≈ 40 ms), and is significantly lower to yet incomplete 70° and 97.6° orbits. We also find evidence of Starlink inter-satellite links (ISLs) delivering superior performance to terrestrial paths for connecting remote regions.

(4) Our fine-grained measurements from terminals in two European countries confirm that Starlink performs network reconfigurations every 15 s, leading to noticeable latency and throughput degradations at sub-second granularity. By correlating data from our terminals in Germany (within 53°) and Scotland (restricted to 70° and 97.6° coverage), we find that the reconfigurations are globally synchronized and likely independent of satellite handovers.

Leveraging multi-dimensional, global, and controlled high resolution measurements, our findings distinctively advance the state-of-the-art by illuminating Starlink’s performance and the influence of internal network operations on real-time applications. To foster reproducibility and enable future research, we publish our > 300 GB collected dataset and associated scripts at [44] and [45].




Figure 2: Starlink follows “bent-pipe” connectivity as traffic traverses the client-side terminal, one or more satellites via inter-sat links (ISLs), nearest ground station (GS), ingressing with the terrestrial Internet via a point-of-presence (PoP).

2 BACKGROUND

Starlink is a LEO satellite network operated by SpaceX that aims to provide global Internet coverage through a fleet of satellites flying at ≈ 500 km above the Earth’s surface. The majority of Starlink’s operational 4000 satellites lie within the 53° shell, which only covers parts of the globe (see Figure 1). The 70° and 97.6° orbits allow serving regions near the poles. These other shells however have fewer satellites (see Appendix A, Table 2 for constellation details).

Figure 2 shows the cross-section of Starlink end-to-end connectivity. To access the Internet over the Starlink network, end-users require a dish, a.k.a. “Dishy”¹, that communicates with satellites visible above 25° of elevation through phased-array antennas using Ku-band (shown as User Link (UL)). Starlink satellites, equipped with multiple antennas subdivided into beams, can connect to multiple terminals simultaneously [18] and relay all connections to a ground station (GS) on a Ka-band link (shown in green). The connection forms a direct “bent-pipe” in case the terminal and GS lie within a single satellite’s coverage cone; otherwise, the satellites can relay within space to reach far-off GSs via laser inter-satellite links (ISLs), forming an “extended bent-pipe”. Note that not all Starlink satellites are ISL-capable and it is difficult to effectively estimate ISL usage as Starlink satellites have no visibility at IP layer [52]. Finally, the GSs terrestrially relay traffic to Starlink points-of-presence (PoPs), which route it to the destination server via terrestrial Internet [6]. No official source exists, so we rely on crowd-sourced data for the geolocations of GSs and PoPs [51], which is also shown in Figure 1.

3 MEASUREMENT METHODOLOGY

3.1 Global Measurements

Measurement Lab (M-Lab) M-Lab [12] is an open-source project that allows users to perform end-to-end throughput and latency speed tests from their devices to 500+ servers in 60+ metropolitan areas [30]. Google offers M-Lab measurements when a user searches for “speed test” [29], serving as the primary source of measurement initiations [12, 20, 54]. At its core, M-Lab uses the Network Diagnostic Tool (NDT) [39], which measures uplink and downlink performance using a single 10 s WebSocket TCP connection. The platform also records fine-grained transport-level metrics (`tcp_info`), including goodput, round-trip time (RTT) and losses, along with IP, Autonomous System Number (ASN), and geolocation of both the end-user device and the selected M-Lab server.

¹We use “Dishy” and “user terminal” interchangeably in the paper.




Figure 3: Overview of global Starlink measurements in this study. Heatmap denotes M-Lab speedtest measurement densities. Starlink RIPE Atlas probes are shown as red circles.

We identify measurements from the Starlink clients via their ASN (AS14593). The M-Lab dataset includes samples from 59 out of 63 countries where Starlink is operational. We restrict our analysis to `ndt7` measurements, which use TCP BBR and countries with *at least 1000 measurements* since June 2021 (launch of Starlink v1.0 satellites [28]), resulting in 19.2 M M-Lab measurement samples from 34 countries. M-Lab infers the approximate location of Starlink users from their public IP which is assigned by the PoP [69]. As a result, all speed tests across countries are mapped to a city, except for USA and Canada, which are sub-divided into multiple regions. While we examine such artifacts by contrasting the M-Lab and RIPE Atlas results (§6.1), we approached our analysis with caution, particularly when examining fine-grained region-specific insights.

RIPE Atlas. RIPE Atlas is a measurement platform that the networking research community commonly employs for conducting measurements [59]. We utilized 98 Starlink RIPE Atlas probes across 21 countries (see Figure 3). Our measurement targets were 145 data centers from *seven* major cloud providers – Amazon EC2, Google, Microsoft, Digital Ocean, Alibaba, Amazon Lightsail, and Oracle (see Appendix B). The chosen operators represent the global cloud market [3, 24, 32, 78] and ensure that our endpoints are close to Starlink PoPs, which are usually co-located with Internet eXchange Point (IXP) or data center facilities [19, 25]. We perform ICMP traceroutes from Atlas probes to endpoints situated on the same or neighboring continent. We extract and track per-hop latencies between Starlink probe terminal-to-GS (identified by static `100.64.0.1` address), GS-to-PoP (`172.16/12`) and PoP-to-endpoint at 2 s intervals [52]. Additionally, we also extract semantic location embeddings in reverse DNS PTR entry, e.g. `tata-level3-seattle2.level3.net` to further improve geolocation accuracy [34]. Our experiments over *ten* months (Dec 2022 to Sept 2023) resulted in ≈ 1.8 M measurement samples.

3.2 Real-time Web Application Measurements

Zoom Video Conferencing. We set up a Zoom call between our server, with access to an unobstructed Starlink dish and 1 Gbps Ethernet, and an AWS machine located close to the assigned Starlink PoP. We set up virtual cameras and microphones on both machines, which were fed by a pre-recorded video of a person talking, resulting in bidirectional transmission. Both servers were synchronized to local stratum-1 NTP servers and we recorded (and analyzed) sub-second Zoom QoS metrics using the toolchain from [41].

Cloud Gaming. We leverage the automated system from [17] to evaluate the performance of playing the racing game “The Crew” on the Amazon Luna [2] platform. The system uses (i) a custom streaming client to record end-to-end information about media




Figure 4: Field-of-view experiment setup. Dishy, deployed at a high latitude location, is obstructed by a metal shielding, which restricts its connectivity to the 70° and 97.6° orbits.

streams, such as frame and bitrate, and (ii) a bot that executes in-game actions at pre-defined intervals. We compute the *game delay* as the time passed since the input action was triggered in post-processing. Amazon Luna serves games at 60 FPS and 1920×1080 resolution which adaptively reduces to 1280×720 depending on the network. We ran the game streaming client on the same machine as the Zoom measurements, additionally setting up a 5G modem to compare Starlink against cellular. Similar to Zoom, the Luna game server was on AWS server close to our Starlink PoP (≈ 1 ms RTT).

3.3 Targeted Measurements

A significant limitation of RIPE Atlas measurements is their lack of sub-second visibility, which is essential for understanding the intricacies of Starlink network. To combat this, we orchestrated a set of precise, tailored, and controlled experiments, utilizing two Starlink terminals as vantage points (VPs) situated in Germany and Scotland. Our dish in Germany connects to the 53° shell while Scotland dish, due to the high latitude location, can be shielded to confine its communication to the 70° and 97.6° orbits. We do this by placing a metal sheeting as Faraday shield barrier at the South-facing angle of the terminal, which obstructs its view from the 53° inclinations (see Figure 4). We verify with external satellite trackers [27, 57] that the terminal only received connectivity from satellites in 97° or 70° inclinations, which resulted in brief *connectivity windows* followed by periods of no service. We performed experiments using the Isochronous Round-Trip Tester (`irtt`) [55] and `iperf` [16] tools. The `irtt` setup records RTTs at high resolutions (3 ms interval) by transmitting small UDP packets. The `irtt` servers were deployed on cloud VMs in close proximity to the assigned Starlink PoP of both VPs (within 1 ms) – minimizing the influence of terrestrial path on our measurements. We used `iperf` to measure both uplink and downlink throughput and record performance at 100 ms granularity. Simultaneously, we polled the gRPC service on each terminal [65] every second to obtain the connection status information.

4 GLOBAL STARLINK PERFORMANCE

We use the minimum RTT (`minRTT`) reported during `ndt7` tests to the closest M-Lab server globally as baseline since it is not affected by queuing delays due to bandwidth-capped data transfers during speed tests. For context, we also select measurements from devices connected to the top-3 mobile network operators (MNOs) in each country [43]. Note that our filtration results in a mix of wired and wireless access networks since M-Lab does not provide a way to distinguish between the two. We use the same endpoints for both Starlink and terrestrial networks (see §3.1).




Figure 5: Median of minimum RTT (in ms) of devices connected via Starlink (left) and top-3 serving ISPs (right) in the same country to the nearest M-Lab server.





Figure 7: Distributions of M-Lab minRTTs over Starlink from select cities in Europe and South America, respectively.

Global View. Figure 5 shows that, for a majority of countries, clients using terrestrial ISPs achieve better latencies over Starlink. While the median latency of Starlink hovers around 40–50 ms, this distribution varies significantly across geographical regions. For instance, in Colombia, Starlink reports better latencies than established terrestrial networks. Conversely, in Manila (The Philippines), Starlink’s performance is notably inferior (Figure 6). The uneven distribution of GSs and PoPs (Figure 1) may explain the latency differences; the USA, which experiences significantly lower latencies, also boasts a robust ground infrastructure. We observe similar trends in Kenya and Mozambique where the closest PoP is located in Nigeria.

Well-Provisioned Regions. Even though a significant portion of global Starlink measurement samples originate from Seattle ($\approx 10\%$), the region shows consistently low latencies, with the 75th percentile well below 50 ms (Figure 6). Contributing factors can be dense GS availability or service prioritization for Starlink’s headquarters. We also observe that Starlink performance is fairly consistent across the USA, confirming that Seattle is not an anomaly but the norm (see Figure 22a in Appendix C). This highlights the LEO network’s potential to bridge Internet access disparities, which currently affects the quality of terrestrial Internet in the USA [37, 53]. Europe is also relatively well covered with GSs but hosts only three PoPs that are in the UK, Germany, and Spain. Proximity to the nearest PoP correlates strongly with minRTT performance in Figure 7 – Dublin, London, and Berlin exhibit comparable latencies to the US, while for Rome and Paris, the 75th percentile is ≈ 20 ms longer. Unlike US, Starlink latencies in EU has longer tail, often surpassing 100 ms.

Under-Provisioned Regions. Starlink’s performance in Colombia hints at its potential for connecting under-provisioned regions. However, Figure 7 shows that Starlink in South America (SA) trails significantly behind the US and Europe, with the 75th percentile exceeding 100 ms and tail reaching 200 ms. Oceania also shows similar performance (see Figure 22b in Appendix C). By extracting the share of satellite vs. terrestrial path (i.e PoP \leftrightarrow M-Lab servers, see Figure 18 in Appendix C)², we find that the majority of SA Starlink latency is due to the bent-pipe. In contrast, latencies from Mexico and Africa (except Nigeria) show terrestrial influence, which we allude to non-optimal PoP assignments by Starlink routing.

²We subtract the latency to the Starlink PoP reported by M-Lab’s reverse traceroutes from the end-to-end TCP minRTT.




Figure 8: RTT inflation (maxRTT-minRTT) during M-Lab Starlink speedtests per continent: (a) download, (b) upload traffic.

We also observe an interesting impact of ground infrastructure in the Philippines, where a local PoP was deployed in May 2023. Prior to this, Starlink traffic from the country was directed to the nearest Japanese PoP, traversing long submarine links to circle back to the geographically closest M-Lab server in Philippines – evident from additional 50–70 ms RTT in Figure 19 in Appendix C for Philippines users to reach in-country vs. Japanese M-Lab servers. However, post-May 2023, the latencies to in-country servers reduced by 90% as the traffic was routed via the local PoP. Despite such artifacts, Starlink shows an evident trend towards more consistent sub-50 ms latencies globally over the past 17 months, specifically evident in Sydney (Figure 21a in Appendix C). We conclude that while Starlink slightly lags behind terrestrial networks today, the gap will continue to shrink as the ground (and satellite) infrastructure expands.

Latency Under Load. Recent findings suggest that Starlink may be susceptible to *bufferbloat* [13, 23, 48], wherein latencies during traffic load can increase significantly due to excessive queue buildups [40]. To explore this globally, we evaluate the RTT inflation, i.e., the difference between the maximum and minimum RTT observed during a speed test. Figure 8 reveals significant delay inflation under load as during active downloads, Starlink experiences ≈ 2 – $4\times$ increased RTTs, reaching almost 400–500 ms (Figure 8a). While such inflations are consistent across *all* Starlink service areas, they are more prominent in regions with subpar baseline performance, e.g., Mexico. Note that the Starlink latency under load is not symmetric. The 60th percentile of RTT during uploads increases to ≤ 100 ms globally (see Figure 8(b)) compared to ≈ 200 ms during downloads. We observe similar behavior while conducting *iperf* over our controlled terminals. Possible explanations can be queue size differences at the Dishy (affecting uploads), the ground station (affecting downloads), or satellites (impacting both). It is also plausible that Starlink employs active queue management (AQM) techniques [1] to moderate uplink latencies under congestion. This approach, however, may adversely impact applications that demand both high bandwidth and low latency – which we explore in §5.

Goodput. Figure 9 shows Starlink download and upload goodputs from speed tests globally. Unlike latencies (Figure 6), the goodput distributions appear relatively homogeneous. Most Starlink clients




Figure 9: Distribution of median (a) download and (b) upload goodput over Starlink from selected cities globally.




Figure 10: Uplink Zoom traffic over a terrestrial (left) and Starlink (right). Vertical lines show 15 s reconfigurations.

achieve $\approx 50\text{--}100$ Mbps download and $\approx 4\text{--}12$ Mbps upload rates at the 75th percentile. We also do not find any correlation between baseline latencies (see Figure 6) and upload/download goodput, evident from the contrasting cases of Dublin and Manila. However, we observe an inverse correlation between loss rates and goodputs; increasing from 4–8% at the 75th-percentile (see Figure 22 in Appendix C). Seattle, notable for its latency performance, records average goodputs. Considering high measurement density from this region, the trend might be due to Starlink’s internal throttling or load-balancing to prevent congestion [68]. We also find that over the past 17 months, Starlink goodputs have stabilized rather than increased, with almost all geographical regions demonstrating similar performance (shown in Figure 21 in Appendix C).

Takeaway #1 – Starlink exhibits competitive performance to terrestrial ISPs on a global scale, especially in regions with dense GS and PoP deployment. However, noticeable degradation is observable in regions with limited ground infrastructure. Our results further confirm that Starlink is affected by bufferbloat. Starlink appears to be optimizing for consistent global performance, albeit with a slight reduction in goodput.

5 REAL-TIME APPLICATION PERFORMANCE

While the global Starlink performance in §4 is promising for supporting web-based applications, it does not accurately capture the potential impact of minute network changes caused by routing, satellite switches, bufferbloating, etc., on application performance. Real-time web applications are known to be sensitive to such fluctuations [7, 17, 40]. In this section, we examine the performance of Zoom and Amazon Luna cloud gaming over Starlink (see §3.2 for measurements details). This allows us to assess the suitability of the LEO network to meet the requirements of the majority of real-time Internet-based applications, as both applications impose a strict latency control loop. Cloud gaming necessitates high downlink bandwidth, while Zoom utilizes uplink and downlink capacity.

Zoom Video Conferencing. Figure 10 shows samples from Zoom calls conducted over a high-speed terrestrial network and over Starlink. The total uplink throughput over Starlink is slightly higher,

	Terrestrial	Cellular	Starlink
Idle RTT (ms)	9	46	40
Throughput (Mbps)	1000	150	220
Frames-per-second	59±1.51	59±1.68	59±1.63
Bitrate (Mbps)	23.08±0.38	22.82±4.24	22.81±2.16
Time at 1080p (%)	100	94.11	99.45
Freezes (ms/min)	0±0	0±220.34	0±119.74
Inter-frame (ms)	17±3.65	18±11.1	16±6.76
Game delay (ms)	133.53±19.79	165.82±23.55	167.13±23.12
RTT (ms)	11±13.41	39±17.06	50±16.28
Jitter buffer (ms)	15±3.27	12±1.33	15±3.35

Table 1: Luna gaming results over 150 mins playtime. Values denote median±SD and the worst performer is highlighted.




Figure 11: Cloud gaming over 5G (left) and Starlink (right). Vertical dashed lines show Starlink reconfiguration intervals.

which we trace to FEC (Forward Error Correction) packets that are frequently sent in addition to raw video data (on average 146±99 Kbps vs. 2±2 Kbps over terrestrial). The frame rate, inferred from the packets received by the Zoom peer, does not meaningfully differ between the two networks (≈ 27 FPS). Note that, since Zoom does not saturate the available uplink and downlink capacity, it should not be impacted by bufferbloating. Yet, we observe a slightly higher loss rate over LEO, which the application combats by proactively utilizing FEC. The uplink one-way delay (OWD) over Starlink is higher and more variable compared to the terrestrial connection (on average 52±14 ms vs. 27±7 ms). All observations also apply to the downlink except that Starlink’s downlink latency (35±11 ms) is similar to the terrestrial connection (32±7 ms). Our analysis broadly agrees with [79] but our packet-level insight reveals bitrate fluctuations partly caused by FEC. Further, our Starlink connection was more reliable and we did not experience second-long outages.

Interestingly, we observe that the Starlink OWD often noticeably shifts at interval points that occur at 15 s increments. Further investigation reveals the cause to be the Starlink *reconfiguration interval*, which, as reported in FCC filings [72], is the time-step at which the satellite paths are reallocated to the users. Other recent work also reports periodic link degradations at 15 s boundaries in their experiments, with RTT spikes and packet losses of several orders [25, 52, 74]. We explore the impact of reconfiguration intervals and other Starlink-internal actions on network performance in §6.

Amazon Luna Cloud Gaming. Table 1 shows 150 minutes of cloud gaming performance over terrestrial, 5G cellular, and Starlink networks. Overall, all networks realized close to 60 FPS playback rate at consistently high bitrate (≈ 20 Mbps). Starlink lies in between the better-performing terrestrial and cellular in terms of bitrate




Figure 12: Last-mile latencies for different countries. “Starlink” denotes satellite bent-pipe over RIPE Atlas while “Cellular” wireless access from Speedchecker [10].

fluctuations, frame drops and freezes³. Starlink exhibits the highest game delay, i.e., the delay experienced by the player between issuing a command and witnessing its effect. Specifically, the wired network delivers the visual response about 2 frames (≈ 33 ms) earlier than both 5G and Starlink. While examining the gaming performance over time, we observe occasional drops to < 20 FPS over Starlink (see Figure 11), that coincide with Starlink’s reconfiguration interval. These fluctuations are only visible at sub-second granularity and, hence, are not reflected in global performance analysis (§4).

Despite these variations, Starlink’s performance remains competitive with 5G, highlighting its potential to deliver real-time application support, especially in regions with less mature cellular infrastructure. Note, however, that our Starlink terminal was set up without obstructions and the weather conditions during measurements were favorable to its operation [36]. Different conditions, especially mobility, may change the relative performance of Starlink and cellular, which we plan to explore further in the near future.

Takeaway #2 – Starlink’s performance is competitive with the current 5G deployment for supporting demanding real-time applications. We also observe that Starlink experiences regular performance fluctuations every 15s linked to its reconfiguration intervals. While these internal black-box parameters do influence performance to a certain extent, application-specific corrective measures, like FEC, are effective in mitigating these artifacts.

6 DISSECTING THE BENT-PIPE

We now attempt to uncover Starlink’s behind-the-scenes operations and their impact on network performance. We follow a two-pronged approach to undertake this challenge. Our longitudinal traceroute measurements over RIPE Atlas accurately isolate the bent-pipe (terminal-to-PoP) global performance, allowing us to correlate it with parameters like ground station deployment, satellite availability, etc. (§6.1). We then perform controlled, high-resolution experiments over Starlink terminals deployed in two EU countries to zoom in on bent-pipe operation and highlight traffic engineering signatures that may impact application performance (§6.2).

6.1 Global Bent-Pipe Performance

Starlink vs. Cellular Last-mile We contrast our end-to-end M-Lab and real-time application analysis by comparing the Starlink bent-pipe latencies from RIPE Atlas traceroutes to cellular wireless last-mile (device-to-ISP network) access. Given the under-representation of cellular probes in RIPE Atlas, we augment our dataset with recent measurements from Dang et al. [10], which

³Freeze is when the inter-frame delay (IFD) is larger than $\max(3 \times \text{IFD}, \text{IFD} + 150)$.




Figure 13: Correlation between Starlink bent-pipe latency and Dishy-GS distance. Red line denotes linear regression fit.

leveraged 115,000 cellular devices worldwide over the Speedchecker platform. Figure 12 presents a comparative analysis of Starlink and cellular last-mile across countries common in both datasets. Consistent with our previous findings, we find that the Starlink bent-pipe latencies fall within 36–48 ms, with the median hovering around 40 ms for almost all countries. Similarly, we find consistent cellular last-mile latencies across all countries, but almost 1.5 \times less than Starlink. Recent investigations [42] report similar access latencies over WiFi and cellular networks. The bent-pipe latencies also corroborate our estimations in §4 that the terminal \leftrightarrow PoP path is the dominant contributor to the end-to-end latency. Out of the 21 countries with Starlink-enabled RIPE Atlas probes, the only exceptions where the bent-pipe latency is significantly higher (≈ 100 ms) are the Virgin Islands (US), Reunion Islands (FR), and Falkland Islands (UK). Correlating with Figure 3, we find that Starlink neither has a GS nor a PoP in these regions, which may result in traffic routing over ISLs to far-off GS leading to longer bent-pipe latencies.

Impact of Ground Infrastructure. We extend our analysis by exploring the correlation between the distance from Starlink users to the GS and bent-pipe latencies. Recall that we rely on crowd-sourced data [51] for geolocating Starlink ground infrastructure since these are not officially publicly disclosed. We deduce through our traceroutes that Starlink directs its subscribers to the nearest GS relative to the PoP, as the GS \leftrightarrow PoP latencies are ≈ 5 ms (almost) globally (see Figure 20 in Appendix C – sole exceptions being US and Canada with 7–8 ms, likely due to abundant availability of GSs and PoPs causing routing complexities). Figure 13 correlates the reported bent-pipe latency with the terminal \leftrightarrow GS distance. Each point in the plot denotes at least 1000 measurements. We observe that bent-pipe latencies tend to increase with increasing distance to the GS. Furthermore, we find that the predominant distance between GS and the user terminal is ≤ 1200 km, which is also the approximate coverage area width of a single satellite from 500 km altitude [5] – suggesting that these connections are likely using direct bent-pipe, either without or with short ISL paths. Few terminals, specifically in Reunion, Falkland and the Virgin Islands, connect to GSs significantly farther away, possible only via long ISL chains, which we analyze further as a case study.

Case Study: Reunion Island. The majority of Starlink satellites (starting from v1.5 deployed in 2021) are equipped with ISLs [75], and reports from SpaceX suggest active utilization of these links [76]. Recent studies also agree with the use of ISLs [19], but point out inefficiencies in space routing [33, 74]. Nonetheless, the invisibility of satellite hops in traceroutes poses a challenge in accurately assessing the use or impact of ISLs. As such, we focus on a probe in Reunion Island (RU), which connects to the Internet via Frankfurt




Figure 14: Bent-pipe RTT segments from Reunion Island (yellow) vs. Germany (red) connecting to Germany PoP. Vertical lines show latency over Atlas probes connected via fiber from both locations to the Frankfurt server (PoP location).

PoP (≈ 9000 km). Figure 14 segments the bent-pipe RTT between the user terminal (Dishy) to GS (non-terrestrial), and from GS to the PoP (terrestrial). For comparison, we also plot the RTTs from a probe within Germany (DE) connecting to the same PoP (≈ 500 km, in red). The vertical lines represent the median RTT over terrestrial infrastructure from both probe locations to the PoP. Firstly, we observe minimal GS \leftrightarrow PoP latency for both locations, verifying that the RU satellite link is using ISLs. Secondly, in RU, Starlink shows significant latency improvement over fiber (≈ 60 ms). This is because the island has limited connectivity with two submarine cables routing traffic 10,000 km away, either in Asia or South America [49]. Starlink provides a better option by avoiding the terrestrial route altogether, directly connecting RU users to the dense backbone infrastructure in EU [8]. However, since the bent-pipe incurs at least 30–40 ms latency in the best-case, Starlink is less attractive in regions with robust terrestrial network infrastructure (also evident from the DE probe where fiber achieves better latencies).

Impact of Serving Orbit. Recall that the majority of Starlink satellites are deployed in the 53° inclination (see Table 2 in Appendix A). Consequently, network performance for clients located outside this orbit’s range may vary widely as they are serviced by few satellites in 70° and 97.6° orbits. Figure 15 contrasts the bent-pipe latencies of probe in Alaska (61.5685N , 149.0125W) [“A”] to probes within 53° orbit. Despite dense GS availability, the bent-pipe latencies for Alaska are significantly higher ($\approx 2\times$). The Swedish probe [“B”] at 59.6395N is at the boundary of 53° orbit but still exhibits comparable latency to Canada, UK, and Germany. Furthermore, the Alaskan probe experiences intermittent connectivity, attributed to the infrequent passing of satellite clusters within the 70° and 97.6° orbits. These findings indicate substantial discrepancies in Starlink’s performance across geographical regions, which may evolve for the better as more satellites are launched in these orbits. Nevertheless, we leverage the sparse availability of satellites at higher latitudes to further dissect the bent-pipe operations in §6.2.

Takeaway #3 – The Starlink “bent-pipe” accounts for ≈ 40 ms latency globally. In certain cases with ISLs use, the latencies might escalate, yet still outshine traditional terrestrial networks when bridging remote regions. The satellite link yields stable latencies, provided that the client is served by the dense 53° orbit.

6.2 Controlled Experiments

Global Scheduling. We performed simultaneous iRTT measurements from terminals in Edinburgh (UK) and Munich (DE). Note that the countries are sufficiently geographically removed that




Figure 15: Bent-pipe latencies for “A” (in Alaska) covered by the 70° and 97.6° while the rest (Sweden “B”, Canada “C”, UK “D”, and Germany “E”) are also covered by 53° .

both cannot be connected to the same serving satellite and are assigned different PoPs. The resulting RTTs, shown in Figure 16a, vary consistently, being comparatively stable within each Starlink reconfiguration interval but potentially changing between intervals. Moreover, the time-wise alignment of reconfiguration intervals for both vantage points indicates that Starlink operates on a globally coordinated schedule, rather than on a per-Dishy or per-satellite basis. These results are in line with other recent studies [74], which also hint that Starlink utilizes a global network controller. Previous studies [11] have noticed drops in downlink throughput every 15 s but have not correlated these with the reconfiguration intervals. We also observe throughput drops on both downlink and uplink, shown in Figure 16b, that occur at the reconfiguration interval boundaries. Similar to the RTT, the throughput typically remains relatively consistent within an interval, but fluctuates between intervals. These results corroborate the periodic degradations observed in §5.

Disproving Satellite Handoff Hypothesis. Previous works have suggested satellite or beam changes at reconfiguration interval boundaries to be the root-cause of network degradation [11, 64, 74]. To investigate this hypothesis, we deliberately obstructed the field-of-view of our UK terminal to prevent it from connecting to the dense 53° shell (see §3.3 for details). The restriction curtailed the number of candidate (potentially connectable) satellites to 13%, resulting in intermittent connectivity. By synchronizing the timings of each connectivity window with the overhead positions of candidate satellites (from CelesTrak [27] and other sources [57]), we identify several windows where the terminal can be served by only a single satellite. Figure 16c (upper) shows RTTs from such a window. The significant RTT variance between intervals invalidates the hypothesis that the RTT changes are caused by satellite handovers (no handoffs are possible with single satellite in field-of-view). Similar to RTT, we also witness uplink and downlink throughput drops at interval boundaries even when single candidate satellite is visible.

Scheduling Updates. Figure 16c (lower) shows the distribution of start and end times of the connectivity windows during our restricted field-of-view experiments. We observed a strong correlation between connectivity end times and reconfiguration interval (RI) boundary, which is not seen with start times⁴. The result hints at internal network scheduling changes at reconfiguration interval boundaries, i.e., Starlink assigns its terminals new satellites (or frequencies) every 15s. We hypothesize that with an obstructed view, the scheduler cannot find better alternatives in the 70° and 97.6° orbits, resulting in connectivity loss at the end of the window.

⁴The fact that many appear to end 1s after the boundary is an artifact of the limited (per-second) granularity of the gRPC data and that the gRPC timestamps originate from the client making the gRPC requests rather than the user terminal.




Figure 16: (left, a) iRTT latencies with Dishys in two countries connected to different ground infrastructure; (middle, b) Maximum uplink and downlink throughput over a 195-second (13 interval) period; (right, c) (upper) RTTs for a connectivity window where the Dishy was connected to only a single satellite; (lower) Probability distribution of the time between the connectivity window start / end and the previous reconfiguration interval (RI). Vertical dashed lines show Starlink reconfiguration intervals.

Analysis Summary. Putting together our various observations, we theorize that Starlink relies on a global scheduler that re-allocates the user \leftrightarrow satellite(s) \leftrightarrow GS path every 15s. An FCC filing from Starlink implies this behavior [64] and recent studies also suggest that the LEO operator performs periodic load balancing at reconfiguration boundaries, reconnecting all active clients to satellites [19, 74]. The theory also explains our observed RTT and throughput changes when only a single candidate satellite is in view. It is plausible that Starlink may have rescheduled the terminal to the same satellite but with reallocated frequency and routing resources. Regardless, these reconfigurations result in brief sub-second connection disruptions, which may become more noticeable at the application-layer as the number of subscribers on the network increases over time.

Takeaway #4 – Starlink uses 15s-long reconfiguration intervals to globally schedule and manage the network. Such intervals cause latency/throughput variations at the interval boundaries. Hand-offs between satellites are not the cause of these effects. Indeed, our findings hint at a scheduling system reallocating resources for connections once every reconfiguration interval.

7 RELATED WORK

LEO satellites have become a subject of extensive research in recent years, with a particular focus on advancing the performance of various systems and technologies. Starlink, the posterchild of LEO networks, continues to grow in its maturity and reach with > 2M subscribers as of September 2023 [60]. Despite its growing popularity, there has been limited exploration into measuring Starlink’s performance so far. Existing studies either have a narrow scope, employing only a few vantage points [11, 35, 40] or focus on broad application-level operation [25, 79] without investigating root-causes. Few studies have looked into the mobile behavior of Starlink and compared it to terrestrial cellular carriers [14, 36].

A few endeavors have attempted to unveil the operations of Starlink’s black-box network. Pan et al. [52] revealed the operator’s internal network topology from traceroutes, whereas Tanveer et al. [74] spotlighted a potential global network controller. The absence of global measurement sites poses a predominant challenge hampering a comprehensive understanding of Starlink’s performance. As we show in this work, Starlink’s performance varies

geographically due to differing internal configurations and ground infrastructure availability. Some researchers have devised innovative methods to combat this. For example, Izhikevich et al. [19] conducted measurements towards exposed services behind the Starlink user terminal, while Taneja et al. [73] mined social media platforms like Reddit to gauge the LEO network’s performance. Our study not only corroborates and extends existing findings but also stands as the most extensive examination to date. Our approach – anchored in detailed insights from 34 countries, leveraging 19.2 million crowdsourced M-Lab measurements, 1.8 million active RIPE Atlas measurements, and two controlled terminals connecting to different Starlink orbits – provides a deeper understanding of the Starlink “bent-pipe” and overall performance.

8 CONCLUSIONS

Despite its potential as a “global ISP” capable of challenging the state of global Internet connectivity, there have been limited performance evaluations of Starlink to date. We conducted a multi-faceted investigation of Starlink, providing insights from a global perspective down to internal network operations. Globally, our analysis showed that Starlink is comparable to cellular for supporting real-time applications (in our case Zoom and Luna cloud gaming), though this varies based on proximity to ground infrastructure. Our case study shows Starlink inter-satellite connections helping remote users achieve better Internet service than terrestrial networks. However, at sub-second granularity, Starlink exhibits performance variations, likely due to periodic internal network reconfigurations at 15s intervals. We find that the reconfigurations are synchronized globally and are not caused by satellite handovers. As such, this first-of-its-kind study is a step towards a clearer understanding of Starlink’s operations and performance as it continues to evolve.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their valuable feedback. This work was supported by 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 also thank the Starlink mailing list [71] which motivated several investigations within this work. Finally, we thank the Measurement Lab and RIPE Atlas platforms for providing us access to their dataset and infrastructure.

REFERENCES

- [1] Richelle Adams. 2012. Active queue management: A survey. *IEEE communications surveys & tutorials* 15, 3 (2012), 1425–1476.
- [2] Amazon. 2023. Luna Cloud Gaming. <https://luna.amazon.com/>.
- [3] Todd Arnold, Ege Gürmeriçli, Georgia Essig, Arpit Gupta, Matt Calder, Vasileios Giotsas, and Ethan Katz-Bassett. 2020. (How Much) Does a Private WAN Improve Cloud Performance?. In *IEEE INFOCOM 2020 - IEEE Conference on Computer Communications*. 79–88. <https://doi.org/10.1109/INFOCOM41043.2020.9155428>
- [4] Alan Boyle. 2019. Amazon to offer broadband access from orbit with 3,236-satellite 'Project Kuiper' constellation. <https://www.geekwire.com/2019/amazon-project-kuiper-broadband-satellite/>. Accessed: 2023-05-26.
- [5] Shkelzen Cakaj. 2021. The parameters comparison of the "Starlink" LEO satellites constellation for different orbital shells. *Frontiers in Communications and Networks* 2 (2021), 643095.
- [6] Case No. 2021-00002. 2021. *Application of Starlink Services*. https://psc.ky.gov/pseccf/2021-00002/kerry.ingle%40dinslaw.com/01042021010318/Application_-_Designation_as_ETC.PDF Accessed on 2023-05-24.
- [7] Hyunseok Chang, Matteo Varvello, Fang Hao, and Sarit Mukherjee. 2021. Can You See Me Now? A Measurement Study of Zoom, Webex, and Meet. In *Proceedings of the 21st ACM Internet Measurement Conference (Virtual Event) (IMC '21)*. Association for Computing Machinery, New York, NY, USA, 216–228. <https://doi.org/10.1145/3487552.3487847>
- [8] Yi-Ching Chiu, Brandon Schlinker, Abhishek Balaji Radhakrishnan, Ethan Katz-Bassett, and Ramesh Govindan. 2015. Are We One Hop Away from a Better Internet?. In *Proceedings of the 2015 Internet Measurement Conference (Tokyo, Japan) (IMC '15)*. Association for Computing Machinery, New York, NY, USA, 7 pages. <https://doi.org/10.1145/2815675.2815719>
- [9] 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>
- [10] 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>
- [11] Johan Garcia, Simon Sundberg, Giuseppe Caso, and Anna Brunstrom. 2023. Multi-Timescale Evaluation of Starlink Throughput. In *Proceedings of the 1st ACM Workshop on LEO Networking and Communication (Madrid, Spain) (LEO-NET '23)*. Association for Computing Machinery, New York, NY, USA, 31–36. <https://doi.org/10.1145/3614204.3616108>
- [12] Phillipa Gill, Christophe Diot, Lai Yi Ohlsen, Matt Mathis, and Stephen Soltész. 2022. M-Lab: User Initiated Internet Data for the Research Community. *SIGCOMM Comput. Commun. Rev.* 52, 1 (mar 2022), 34–37. <https://doi.org/10.1145/3523230.3523236>
- [13] Carlo Augusto Grazia, Natale Patriciello, Martin Klapez, and Maurizio Casoni. 2017. Mitigating congestion and bufferbloat on satellite networks through a rate-based AQM. In *2017 IEEE International Conference on Communications (ICC)*. 1–6. <https://doi.org/10.1109/ICC.2017.7996776>
- [14] 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>
- [15] Business Insider. 2023. *Everything we know about Elon Musk's Starlink satellites and future internet plans*. <https://www.businessinsider.com/elon-musk-starlink-satellites-internet> Accessed: 2023-10-12.
- [16] iperf. 2023. iperf / iperf3. <https://iperf.fr/>.
- [17] Hassan Iqbal, Ayesha Khalid, and Muhammad Shahzad. 2021. Dissecting Cloud Gaming Performance with DECAF. *Proc. ACM Meas. Anal. Comput. Syst.* 5, 3, Article 31 (dec 2021), 27 pages. <https://doi.org/10.1145/3491043>
- [18] 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.
- [19] Liz Izhikevich, Manda Tran, Katherine Izhikevich, Gautam Akiwate, and Zakir Durumeric. 2024. Democratizing LEO Satellite Network Measurement. <https://arxiv.org/abs/2306.07469>
- [20] Akshath Jain, Deepayan Patra, Peijing Xu, Justine Sherry, and Phillipa Gill. 2022. The Ukrainian Internet under Attack: An NDT Perspective. In *Proceedings of the 22nd ACM Internet Measurement Conference (Nice, France) (IMC '22)*. Association for Computing Machinery, New York, NY, USA, 166–178. <https://doi.org/10.1145/3517745.3561449>
- [21] Amritha Jayanti. 2023. Starlink and the Russia-Ukraine War: A Case of Commercial Technology and Public Purpose? <https://www.belfercenter.org/publication/starlink-and-russia-ukraine-war-case-commercial-technology-and-public-purpose>.
- [22] Yubing Jian, Mohit Agarwal, Shyam Krishnan Venkateswaran, Yuchen Liu, Douglas M. Blough, and Raghupathy Sivakumar. 2020. WiMove: Toward Infrastructure Mobility in MmWave WiFi. In *Proceedings of the 18th ACM Symposium on Mobility Management and Wireless Access (Alicante, Spain) (MobiWac '20)*. Association for Computing Machinery, New York, NY, USA, 11–20. <https://doi.org/10.1145/3416012.3424625>
- [23] Haiqing Jiang, Zeyu Liu, Yaogong Wang, Kyunghan Lee, and Injong Rhee. 2012. Understanding Bufferbloat in Cellular Networks. In *Proceedings of the 2012 ACM SIGCOMM Workshop on Cellular Networks: Operations, Challenges, and Future Design (Helsinki, Finland) (CellNet '12)*. Association for Computing Machinery, New York, NY, USA, 1–6. <https://doi.org/10.1145/2342468.2342470>
- [24] Yuchen Jin, Sundararajan Renganathan, Ganesh Ananthanarayanan, Junchen Jiang, Venkata N. Padmanabhan, Manuel Schroder, Matt Calder, and Arvind Krishnamurthy. 2019. Zooming in on Wide-Area Latencies to a Global Cloud Provider. In *Proceedings of the ACM Special Interest Group on Data Communication (Beijing, China) (SIGCOMM '19)*. Association for Computing Machinery, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3341302.3342073>
- [25] Mohamed M. Kassem, Aravindh Raman, Diego Perino, and Nishanth Sastry. 2022. A Browser-Side View of Starlink Connectivity. In *Proceedings of the 22nd ACM Internet Measurement Conference (Nice, France) (IMC '22)*. Association for Computing Machinery, New York, NY, USA, 151–158. <https://doi.org/10.1145/3517745.3561457>
- [26] Simon Kassing, Debopam Bhattacharjee, André Baptista Águas, Jens Eirik Saethre, and Ankit Singla. 2020. Exploring the "Internet from Space" with Hypatia. In *Proceedings of the ACM Internet Measurement Conference (Virtual Event, USA) (IMC '20)*. Association for Computing Machinery, New York, NY, USA, 214–229. <https://doi.org/10.1145/3419394.3423635>
- [27] Dr. T.S. Kelso. 2023. *CelesTrak*. <https://celestrak.org/> Accessed: 2023-10-12.
- [28] Gunther Krebs. 2021. Starlink Block v1.0. https://space.skyrocket.de/doc_sdat/starlink-v1-0.htm.
- [29] Measurement Lab. 2023. M-Lab to the Cloud. <https://www.measurementlab.net/blog/virtual-sites-gcp>.
- [30] Measurement Lab. 2023. Measurement Lab Infrastructure. <https://www.measurementlab.net/status/>.
- [31] Zeqi Lai, Hewu Li, Yangtao Deng, Qian Wu, Jun Liu, Yuanjie Li, Jihao Li, Lixin Liu, Weisen Liu, and Jianping Wu. 2023. StarryNet: Empowering Researchers to Evaluate Futuristic Integrated Space and Terrestrial Networks. In *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)*. USENIX Association, Boston, MA, 1309–1324. <https://www.usenix.org/conference/nsdi23/presentation/lai-zeqi>
- [32] Ang Li, Xiaowei Yang, Srikanth Kandula, and Ming Zhang. 2010. CloudCmp: Comparing Public Cloud Providers. In *Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement (Melbourne, Australia) (IMC '10)*. Association for Computing Machinery, New York, NY, USA, 14 pages. <https://doi.org/10.1145/1879141.1879143>
- [33] 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>
- [34] Matthew Luckie, Bradley Huffaker, Alexander Marder, Zachary Bischof, Marianne Fletcher, and K Claffy. 2021. Learning to Extract Geographic Information from Internet Router Hostnames. In *Proceedings of the 17th International Conference on Emerging Networking Experiments and Technologies (Virtual Event, Germany) (CoNEXT '21)*. Association for Computing Machinery, New York, NY, USA, 440–453. <https://doi.org/10.1145/3485983.3494869>
- [35] Melisa López, Sebastian Bro Damsgaard, Ignacio Rodriguez, and Preben Mogenssen. 2022-12. An Empirical Analysis of Multi-Connectivity between 5G Terrestrial and LEO Satellite Networks. In *2022 IEEE Globecom Workshops (GC Wkshps)*. IEEE, 1115–1120. <https://doi.org/10.1109/GCWkshps56602.2022.10008752>
- [36] 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]
- [37] Tarun Mangla, Udit Paul, Arpit Gupta, Nicole P. Marwell, and Nick Feamster. 2022. Internet Inequity in Chicago: Adoption, Affordability, and Availability. In *Elsevier SSRN*. <https://doi.org/10.2139/ssrn.4182994>
- [38] Jonathan C. McDowell. 2023. Jonathan's Space Pages: StarlinkStatistics. <https://planet4589.org/space/con/star/stats.html>. Accessed: 2023-10.
- [39] Measurement Lab. (2009-02-11 – 2015-12-21). The M-Lab NDT Data Set. <https://measurementlab.net/tests/ndt>.
- [40] François Michel, Martino Trevisan, Danilo Giordano, and Olivier Bonaventure. 2022. A First Look at Starlink Performance. In *Proceedings of the 22nd ACM Internet Measurement Conference (Nice, France) (IMC '22)*. Association for Computing Machinery, New York, NY, USA, 130–136. <https://doi.org/10.1145/3517745.3561416>

- [41] Oliver Michel, Satadal Sengupta, Hyojoon Kim, Ravi Netravali, and Jennifer Rexford. 2022. Enabling passive measurement of zoom performance in production networks. In *Proceedings of the 22nd ACM Internet Measurement Conference* (Nice, France) (IMC '22). Association for Computing Machinery, New York, NY, USA, 244–260. <https://doi.org/10.1145/3517745.3561414>
- [42] Nitinder Mohan, Lorenzo Corneo, Aleksandr Zavodovski, Suzan Bayhan, Walter Wong, and Jussi Kangasharju. 2020. Pruning Edge Research with Latency Shears. 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, 182–189. <https://doi.org/10.1145/3422604.3425943>
- [43] Nitinder Mohan, Andrew Ferguson, Hendrik Cech, Prakita Rayyan Renatin, Rohan Bose, Mahesh Marina, and Jörg Ott. 2023. Top-3 Mobile Network Operators. https://github.com/nitindermohan/multifaceted-starlink-performance/blob/main/mno_related/mno_list_withasn.csv.
- [44] Nitinder Mohan, Andrew Ferguson, Hendrik Cech, Prakita Rayyan Renatin, Rohan Bose, Mahesh Marina, and Jörg Ott. 2024. *A Multifaceted Look at Starlink Performance: Code*. <https://github.com/nitindermohan/multifaceted-starlink-performance>
- [45] Nitinder Mohan, Andrew Ferguson, Hendrik Cech, Prakita Rayyan Renatin, Rohan Bose, Mahesh Marina, and Jörg Ott. 2024. *A Multifaceted Look at Starlink Performance: Dataset*. <https://doi.org/10.14459/2024mp1734703>
- [46] Arvind Narayanan, Eman Ramadan, Jason Carpenter, Qingxu Liu, Yu Liu, Feng Qian, and Zhi-Li Zhang. 2020. A First Look at Commercial 5G Performance on Smartphones. In *Proceedings of The Web Conference 2020* (Taipei, Taiwan) (WWW '20). Association for Computing Machinery, New York, NY, USA, 894–905. <https://doi.org/10.1145/3366423.3380169>
- [47] Arvind Narayanan, Xumiao Zhang, Ruiyang Zhu, Ahmad Hassan, Shuwei Jin, Xiao Zhu, Xiaoxuan Zhang, Denis Rybkin, Zhengxuan Yang, Zhuoqing Morley Mao, Feng Qian, and Zhi-Li Zhang. 2021. A Variegated Look at 5G in the Wild: Performance, Power, and QoE Implications. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference* (Virtual Event, USA) (SIGCOMM '21). Association for Computing Machinery, New York, NY, USA, 610–625. <https://doi.org/10.1145/3452296.3472923>
- [48] TWIT Tech Podcast Network. 2021. Dave Taht To Elon Musk: 'We are gonna fix bufferbloat'. <http://www.youtube.com/watch?v=c9gLo6Xrwgw>
- [49] Rehan Noordally, Xavier Nicolay, Pascal Anelli, Richard Lorion, and Pierre Ugo Tournoux. 2016. Analysis of Internet Latency: The Reunion Island Case. In *Proceedings of the 12th Asian Internet Engineering Conference* (Bangkok, Thailand) (AINTEC '16). Association for Computing Machinery, New York, NY, USA, 49–56. <https://doi.org/10.1145/3012695.3012702>
- [50] OneWeb. 2023. OneWeb: Connect with Ease. https://oneweb.net/connect_with_ease. Accessed: 2023-05-26.
- [51] Nathan Owens. 2023. Unofficial Crowdsourced Starlink Global Gateways & PoPs. <https://bit.ly/3tykCGW>. Accessed: 2023-05-03.
- [52] Jianping Pan, Jinwei Zhao, and Lin Cai. 2023. Measuring a Low-Earth-Orbit Satellite Network. In *2023 IEEE 34th Annual International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, 1–6. <https://doi.org/10.1109/PIMRC56721.2023.10294034>
- [53] Udit Paul, Jiamo Liu, David Farias-Ilerenas, Vivek Adarsh, Arpit Gupta, and Elizabeth Belding. 2022. Characterizing Internet Access and Quality Inequities in California M-Lab Measurements. In *ACM SIGCAS/SIGCHI Conference on Computing and Sustainable Societies (COMPASS)* (Seattle, WA, USA) (COMPASS '22). Association for Computing Machinery, New York, NY, USA, 257–265. <https://doi.org/10.1145/3530190.3534813>
- [54] Udit Paul, Jiamo Liu, Mengyang Gu, Arpit Gupta, and Elizabeth Belding. 2022. The Importance of Contextualization of Crowdsourced Active Speed Test Measurements. In *Proceedings of the 22nd ACM Internet Measurement Conference* (Nice, France) (IMC '22). Association for Computing Machinery, New York, NY, USA, 274–289. <https://doi.org/10.1145/3517745.3561441>
- [55] Pete Heist. 2023. irtt (Isochronous Round-Trip Tester). <https://github.com/heist/irtt>.
- [56] Stephen R. Pratt, Richard A. Raines, Carl E. Fossa, and Michael A. Temple. 1999. An operational and performance overview of the IRIDIUM low earth orbit satellite system. *IEEE Communications Surveys* 2, 2 (1999), 2–10. <https://doi.org/10.1109/COMST.1999.5340513> Conference Name: IEEE Communications Surveys.
- [57] Mike Puchol. 2023. *Starlink Coverage Tracker*. <https://starlink.sx/> Accessed: 2023-10-12.
- [58] 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>
- [59] RIPE. 2023. RIPE Atlas. <https://atlas.ripe.net/>.
- [60] Via Satellite. 2023. Starlink Surpasses 2 Million Subscribers. <https://bit.ly/3rDop5y>.
- [61] Ankit Singla. 2021. SatNetLab: A Call to Arms for the next Global Internet Testbed. *SIGCOMM Comput. Commun. Rev.* 51, 2 (may 2021), 28–30. <https://doi.org/10.1145/3464994.3465000>

Inclination angle	# Planes	Altitude [km]	# Satellites		
			In Position	Launched	Filed [62, 63]
53°	72	550	1401	1665	1584
53.2°	72	540	1542	1637	1584
70°	36	570	301	408	720
97.6°	10	560	230	230	508

Table 2: Starlink orbital shell design and number of operational satellites as of October 2023 [38].

- [62] SpaceX. 2016. SPACEX NON-GEOSTATIONARY SATELLITE SYSTEM. <https://fcc.report/IBFS/SAT-LOA-20161115-00118/1158350.pdf>.
- [63] SpaceX. 2019. SPACEX NON-GEOSTATIONARY SATELLITE SYSTEM. <https://fcc.report/IBFS/SAT-MOD-20190830-00087/1877671>.
- [64] SpaceX. 2021. PETITION OF STARLINK SERVICES, LLC FOR DESIGNATION AS AN ELIGIBLE TELECOMMUNICATIONS CARRIER. <https://www.fcc.gov/ecfs/search/search-filings/filing/1020316268311>.
- [65] sparky8512. 2023. starlink-grpc-tools. <https://github.com/sparky8512/starlink-grpc-tools>.
- [66] Starlink. 2023. Starlink. <https://www.starlink.com>. Accessed: 2023-05-26.
- [67] Starlink. 2023. *Starlink Direct to Cell*. <https://direct.starlink.com/> Accessed: 2023-10-12.
- [68] Starlink. 2023. Starlink Fair Use Policy. <https://www.starlink.com/legal/documents/DOC-1134-82708-70>.
- [69] Starlink. 2023. Starlink Location IP Mapping. <https://geopip.starlinkisp.net/feed.csv>.
- [70] Starlink. 2023. Starlink maritime. <https://www.starlink.com/maritime>.
- [71] Starlink-Admins. 2023. Starlink Mailing List. <https://lists.bufferbloat.net/pipermail/starlink/>. Accessed: 2024-02-01.
- [72] LLC Starlink Services. 2021. Petition of Starlink Services, LLC for Designation as an Eligible Telecommunications Carrier. Submitted to the Department of Telecommunications and Cable, Commonwealth of Massachusetts.
- [73] Aryan Taneja, Debopam Bhattacharjee, Saikat Guha, and Venkata N. Padmanabhan. 2023. On viewing SpaceX Starlink through the Social Media Lens. arXiv:2307.13441 [cs.NI]
- [74] Hammas Bin Tanveer, Mike Puchol, Rachee Singh, Antonio Bianchi, and Rishab Nithyanand. 2023. Making Sense of Constellations: Methodologies for Understanding Starlink's Scheduling Algorithms. In *Companion of the 19th International Conference on Emerging Networking Experiments and Technologies* (, Paris, France) (CoNEXT 2023). Association for Computing Machinery, New York, NY, USA, 37–43. <https://doi.org/10.1145/3624354.3630586>
- [75] Elon Musk Twitter. 2023. Starlink Inter-Satellite Links. <https://twitter.com/elonmusk/status/1535394359373443073?lang=en>.
- [76] Starlink Twitter. 2023. Starlink Twitter inter-satellite link announcement. <https://twitter.com/Starlink/status/1706718537711337650>.
- [77] Wikipedia. 2023. Starlink. <https://en.wikipedia.org/wiki/Starlink>.
- [78] Bahador Yeganeh, Ramakrishnan Durairajan, Reza Rejaie, and Walter Willinger. 2020. A First Comparative Characterization of Multi-cloud Connectivity in Today's Internet. In *International Conference on Passive and Active Network Measurement*. Springer, 193–210.
- [79] Haoyuan Zhao, Hao Fang, Feng Wang, and Jiangchuan Liu. 2023. Realtime Multimedia Services over Starlink: A Reality Check. In *Proceedings of the 33rd Workshop on Network and Operating System Support for Digital Audio and Video* (Vancouver, BC, Canada) (NOSSDAV '23). Association for Computing Machinery, New York, NY, USA, 43–49. <https://doi.org/10.1145/3592473.3592562>

A STARLINK ORBITAL INFORMATION

Starlink and other emerging LEO satellite constellations, such as OneWeb and Kuiper [4, 50], are termed *megaconstellations* since they combine multiple orbital shells compared to single shell systems like Iridium [56]. Table 2 details the number of satellites and their altitude in Starlink's orbital shells. While discussing Starlink's constellation design, we simplify the orbit into circular orbits.

B DATA CENTER ENDPOINTS

Table 3 shows the distribution of our cloud data center endpoints by cloud provider and deployed continent. Each endpoint is a VM in a compute-capable cloud data center location. Our selection is

	Data centers per continent					
	EU	NA	SA	AS	AF	OC
Amazon EC2 (AMZN)	6	6	1	6	1	1
Google Cloud (GCP)	6	10	1	8	-	1
Microsoft Azure (MSFT)	14	9	1	10	2	3
Digital Ocean (DO)	4	6	-	1	-	-
Alibaba (BABA)	2	1	-	2	-	1
Amazon Lightsail (LTSL)	3	2	-	2	-	1
Oracle (ORCL)	4	4	1	7	-	2
Total	39	38	4	36	3	9

Table 3: Global density of data center endpoints used for RIPE Atlas measurements (§6).

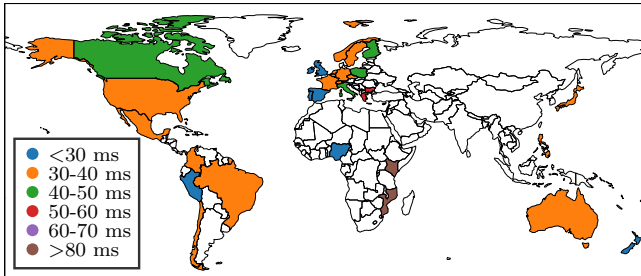


Figure 17: Median Starlink last-mile latencies from M-lab. We subtract the M-lab server ↔ PoP latency in reverse tracroute from overall measured min RTT.




Figure 18: Fraction of satellite link latency. Calculated by dividing latencies in Figure 17 with overall latency.

influenced by previous studies that have found that significant end-to-end performance differences may appear while measuring different cloud networks due to private WANs, peering agreements, etc. [9, 10]. We believe that our comprehensive endpoint selection reduces potential biases due to Internet traffic steering.

C GLOBAL STARLINK PERFORMANCE

Figure 21 provides an overview of Starlink’s performance over one year from selected cities in each continent. We observe decreasing goodputs over time, which can be largely attributed to an increase of Starlink users. The RTT values, while relatively stable, are higher for countries outside of major Starlink operational areas. Sydney (AU) is a notable exception where RTT decreases over time. Figure 19 shows the median minRTT from Philippines to in-country




Figure 19: Comparing median minRTT from tests originating in Manila that targeted M-Lab servers in the Philippines and in Japan. The results are discussed in §4.

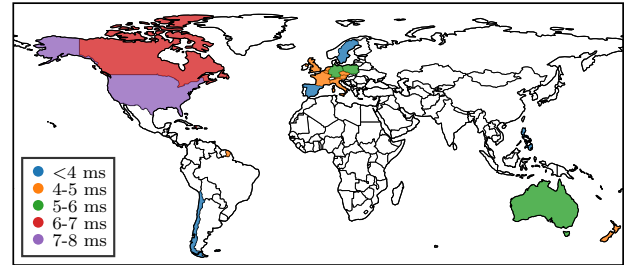


Figure 20: Global GS ↔ PoP latencies from Starlink RIPE Atlas probes.

vs Japanese endpoints. The implications are discussed in §4. Figure 22 shows loss rates during M-Lab download measurements (see §4). Figure 23 complements Figure 7 and shows the distribution of the minimum RTT (minRTT) during M-Lab measurements from selected cities in North America and Oceania. Starlink’s performance in North America varies slightly across different cities and all achieve lower latencies compared to rest of the globe. In Oceania, tests from Auckland and Perth exhibit a similarly low minRTT due to PoPs and GSs near both cities. Latencies in Sydney’s has improved recently (2023/06) as shown in Figure 21a.

D GLOBAL VIEW OF BENT-PIPE OPERATION

Figure 20 augments our discussion on Starlink last-mile performance in §6.1 and shows GS ↔ PoP latencies globally. It is apparent that the latencies are similar all over the world (≤ 6 ms) except in North America (≥ 6 ms), likely due to dense deployment of GSs and PoPs in these regions (see Figure 1).

E TARGETED MEASUREMENT CHALLENGES

We encountered several challenges during our controlled measurements over shielded terminal in §6.2 which we elucidate for researchers who plan to replicate our setup. Specifically, both `irrtt` and `iperf` could not handle interruptions well as they rely on single TCP connection which eventually times out. To overcome this, we replaced `irrtt` with periodic ICMP ping packets sent every 200 ms. On the other hand, we manually controlled `iperf` to overcome connection drops. Specifically, we started `iperf` every time we detected the start of connectivity window (from pings) and stopped the experiment upon interruption. Before starting `iperf` at next window detection, we also restarted the server. Automatically restarting

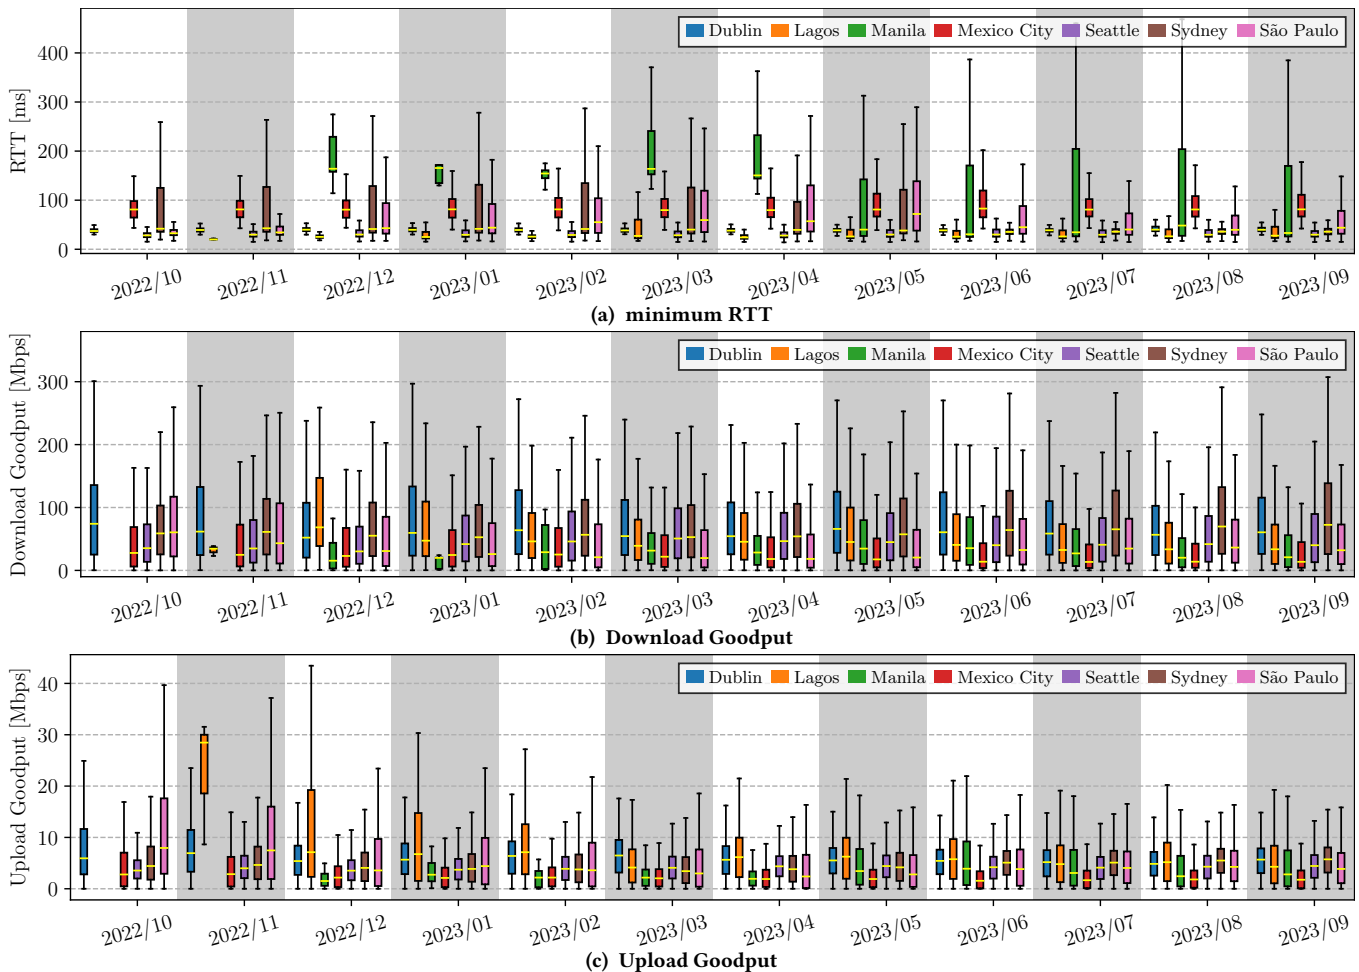


Figure 21: Evolution of Starlink aggregate goodput ((a), (b)) and minimum RTT (c) during download measurements from cities in South America, North America, Europe, and Australia in the last 12 months.




Figure 22: TCP losses during M-Lab downloads from selected cities globally.




Figure 23: The distribution of the minimum RTT (minRTT) during M-Lab measurements from selected cities in North America (a) and Oceania (b).

the iperf server at the end of each connectivity window was not possible because the Starlink-connected computer, now without an Internet connection, could not signal to the remote iperf server.

An additional challenge caused by the interrupted nature of the connection became apparent in the analysis phase. The unstable

connection prevented the clock on the machine connected to the Dishy from synchronising over NTP, resulting in it drifting by several seconds duration of the experiment setup. Accordingly, when the absolute timestamps of the recorded data were analysed, they were adjusted to account for the time slip. The gRPC data was collected from another machine that did not suffer from clock drift.