

Dissecting the Performance of Satellite Network Operators

Aravindh Raman
Telefonica Research

Matteo Varvello
Nokia Bell Labs

Hyunseok Chang
Nokia Bell Labs

Nishanth Sastry
University of Surrey

Yasir Zaki
NYU Abu Dhabi

Abstract

The rapid growth of satellite network operators (SNOs) has revolutionized broadband communications, enabling global connectivity and bridging the digital divide. As these networks expand, it is important to evaluate their performance and efficiency. This paper presents the first comprehensive study of SNOs. We take an opportunistic approach and devise a methodology which allows to identify public network measurements performed via SNOs. We apply this methodology to both M-Lab and RIPE public datasets which allowed us to characterize low level performance and footprint of up to 18 SNOs operating in different orbits. Finally, we identify and recruit paid testers on three popular SNOs (Starlink, HughesNet, and Viasat) to evaluate the performance of popular applications like web browsing and video streaming.

1 Introduction

In recent years, satellite network operators (SNOs) have gained significant attention as an alternative to terrestrial Internet, especially in remote or underserved areas. Companies like Starlink [14], Viasat [18], and HughesNet [3] have invested heavily in satellite technologies to provide high-speed Internet access to their customers regardless of their location. Satellite technologies have evolved from using mainly geosynchronous (GEO) satellites – which operate at around 35,000 km above the earth surface [55] – to Low Earth Orbit (LEO) satellites — which can be as close as 550 km away as in the case of the closest orbital shell from Starlink [26].

Several research papers [48, 53, 56] have investigated the performance of SNOs, especially focusing on Star-

link due to its recent popularity and novel LEO-based technology. The main challenge in performing such measurements is gathering vantage points, *i.e.*, instrumentable devices from where to perform network measurements. To tackle this challenge, researchers have deployed a few controlled nodes with dedicated satellite links (from one to four [56, 53]) or recruited Starlink users to perform network measurements via a browser addon [48].

The main goal of this study is to *scale up* the previous measurements of SNOs in multiple dimensions. More specifically, we aim to expand the scope of the measurements in terms of *space* (*i.e.*, more geographic locations), *time* (*i.e.*, longer measurement periods), *satellite technologies* (*i.e.*, multiple SNOs), and *application diversity* (*i.e.*, multiple applications). In this type of large-scale measurement study, the main challenge is data collection. Indeed, building our own measurement testbed from scratch is challenging, and hard to scale. Not to mention that it would take years to deploy such a testbed, during which lots of data would be lost over time.

Motivated by the above, we opt for an *opportunistic* approach and leverage data and measurement endpoints already available. Our approach is to *identify* SNO measurements from public datasets (M-Lab [8] and Ripe Atlas [11]), and *recruit* SNO users from crowdsourcing platforms (Prolific [10]). Public measurements have the advantage of their scale, *i.e.*, covering many SNOs over a long period of time, but they are limited to low-level measurements, such as **traceroute** and speed tests. Actual SNO users can run more complex measurements, such as Web browsing and video streaming, but likely at a smaller scale in space and time.

For all these different data sources, we first need to correctly identify measurements and testers associated with SNOs. This is challenging since very little ground

*Corresponding author: Aravindh Raman (aravindh.raman@gmail.com)

truth is available, and SNOs are complex entities offering mixed connectivity, e.g., both LEO and GEO as well as satellite and wireline access. To tackle this challenge, we devise the following methodology. First, we leverage ASNs and public information to identify *potential* SNOs along with their access technology (LEO, MEO, or GEO). Next, we use public data from M-Lab to build maps of network latencies per ASN and IP blocks with an AS. We then filter ASNs and IP blocks whose latency profiles show even the slightest incompatibility with the satellite technology offered by their SNO, e.g., less than 500 ms over GEO. This strict filtering identifies 10 SNOs in M-Lab public dataset, and also produces accurate latency profiles per SNO access technology. We then use this information to relax the previous filtering and identify data associated with an additional 8 SNOs, for a total of 18 SNOs.

After applying the above SNO detection mechanism, we identify 11.92 million TCP-based speed tests from 18 SNOs (2 LEO, 1 Medium Earth Orbit or MEO, and 15 GEO) in the M-Lab dataset (between January 2021 and March 2023), and about 6 million `traceroute` measurements from 67 probes connected via Starlink (LEO) in the RIPE Atlas dataset (between May 2022 and May 2023). For Prolific, we find that their prescreening APIs are only partially effective in identifying SNO subscribers. We instead run a *census* which identifies, out of 14,371 participants, 57 potential testers connected via Starlink, HughesNet, and Viasat. Over one month, we recruited 20 out of these 57 testers willing to install and run a browser addon we developed to measure application performance. In the following, we summarize our key findings.

LEO vs. MEO vs. GEO. Our large dataset allows us to perform what we believe is the first large-scale and global characterization of the different SNO technologies available. As expected, we find that LEO supports much higher throughput and much lower latencies than GEO or MEO. However, we also find interesting patterns – for instance, LEO networks suffer from much higher jitter variation, likely due to frequent satellite handoffs. This may negatively affect applications that expect a consistent latency profile. On the other hand, customers of most GEO networks suffer a high amount of data retransmissions, likely due to transport level retransmission timeouts. Fortunately, Performance Enhancing Proxies (PEPs) [69, 72] appear to mitigate this problem and GEO operators which employ PEPs have retransmission profiles similar to LEO operators.

PoP selection matters. On average, Starlink LEO satellites provide connectivity to their customers while only adding an extra 30-40 ms. To achieve such low latency, proper Point-of-Presence (PoP) selection is paramount. We notably identify two main examples of this observation. First, as of today, Starlink customers in the Philippines experience 2x latencies (80 ms) because they are associated to a PoP in Tokyo, Japan (“customer.tkyojpn1.pop.starlinkisp.net”). Second, New Zealand Starlink customers experienced a 20 ms latency reduction since July 2022 when their PoP changed from Sydney (Australia) to Auckland (New Zealand).

Careful technology selection. Modern networking technologies are important for users in GEO-based SNOs like HughesNet and Viasat. Our experiments show that selecting the fastest content-delivery network (CDN), Fastly in our experiments, can reduce download time by up to one second, when loading critical Web component like popular JavaScript (JS) libraries. A similar effect is achieved by JS *minification*, or the process to minimize JS code. Recent evolution of the HTTP protocol, like HTTP/2, also allows to bridge the gap between GEO and LEO users on HTTP/1.1, thanks to features such as connection multiplexing, which result in a reduction in the number of (secure) connections and allow efficient usage of the underlying transport.

2 Background and Related Work

Internet measurement platforms. There are several prominent open platforms for Internet measurements. We here provide more details about the two platforms we leverage in this paper: M-Lab [8] and RIPE Atlas [11]. M-Lab allows researchers to deploy web-based Internet measurement tools to be run by end users in the wild. It operates server pods worldwide, each containing three to four interconnected servers. M-Lab makes the measurement data collected from end users publicly accessible via bulk downloads or Big-Query [7]. RIPE Atlas is a distributed measurement platform powered by thousands of dedicated measurement devices called “probes” hosted by volunteers all around the world. Each probe periodically performs a set of “built-in” measurements: ping, traceroute, DNS, SSL/TLS and HTTP probing. These measurements are primarily directed towards well-known targets, such as DNS root servers and part of the RIPE Atlas infrastructure [12]. Similar to M-Lab, RIPE Atlas makes collected

measurement data publicly available via BigQuery.

Datasets from these platforms have enabled a wide range of measurement studies, e.g., assessing the impact of major societal events [45, 29], characterizing access networks [24, 36], inferring global or regional Internet properties [25, 28], uncovering network anomalies [46, 41], evaluating cloud reachability [32], etc. To the best of our knowledge, however, no prior study has analyzed the data sets with a focus on global SNOs.

Satellite Internet measurements. Kassem et al. [48] studied Starlink connectivity by utilizing a combination of a custom browser extension and dedicated measurement nodes. They observed non-negligible variability in web access performance across different weather conditions and geographic locations, and significant packet loss rates as high as 50%. These findings are based on browser extension data collected from 28 users in 10 cities worldwide and three dedicated nodes provisioned in UK, Spain and USA. The study also compares the performance of traditional access networks (cellular and fiber) with that of LEO-based access. Ma et al. [53] carried out a measurement study on Starlink by utilizing four dedicated Starlink kits stationed over a range of terrains including major cities and remote areas in west Canada. In areas where terrestrial broadband access was available, they also compared Starlink access against terrestrial Internet access. They reported similar observations where Starlink connectivity is subject to more dynamic throughput and latency variation than terrestrial networks and is heavily affected by environments such as terrain characteristics and weather conditions. The authors of [61] performed a detailed study on the traffic carried by a GEO-based SatCom network. By exploiting the fact that all subscriber traffic is relayed via the central ground station, they deployed a traffic monitoring server at the ground station. Their study compares subscriber access patterns across different geographic regions, characterizes RTT, DNS and download performance, and highlights technical challenges faced by the SatCom infrastructure. Finally, several works [31, 34, 33, 56] studied congestion control behaviors and performance of various transport protocols such as TCP, HTTP and QUIC on satellite Internet. These studies relied on a single vantage point, *i.e.*, a controlled lab environment with dedicated satellite links.

Our paper differs from related works in two key aspects. First, it devises a methodology to identify, from public datasets, experiments run via SNOs.

Second, it compares the performance of up to 18 SNOs from multiple different angles (orbit of operations, application, terrestrial connectivity, etc.).

3 Methodology

This section describes the methodology we have devised to dissect the performance of global SNOs. Our approach is to *identify* SNO measurements from public datasets (see Section 3.1), and *recruit* SNO users from crowdsourcing platforms (see Section 3.3). In both approaches, the first step is to correctly identify the measurements or users that are associated with SNOs (see Section 3.2).

3.1 Public Datasets

M-Lab’s Speed Test. This dataset provides upload and download speeds measured by clients distributed across the world using the Network Diagnostic Tool (NDT). Currently, NDT exists in two versions, NDT5 and NDT7. However, the majority of clients have migrated to NDT7 (more than 90% as of 2021 [57]). Accordingly, we focus on the more recent NDT7 traces for our analysis.

An NDT test consists of a single TCP connection between a client and a nearby M-Lab server determined by Google’s location service [38]. The server captures the `TCP_Info` for each speed test experiment running a polling loop on every open TCP socket, thus capturing multiple records per session. The captured trace is uploaded to Google BigQuery, and provides information such as RTT, jitter, delivery rate, and the sizes of sent and received bytes. To overcome a bug in the `APP_Info` values reported by the client [54], we rely on the `TCP_Info` instead. For the latency-based analysis we use 5th percentile value per speed test session. To measure the variability in latencies within a speed test session we measure 95th percentile of jitter and normalize using the 5th percentile latency of that session *i.e.*, jitter variability = $jitter_{p95}/latency_{p5}$. Appendix A discusses the ethics behind our data collection.

RIPE Atlas Built-in Measurements. We gather access to RIPE Atlas between May 3rd, 2022 and May 3rd, 2023. We focus on the built-in measurements (see Section 2) since, in absence of failures, they are deployed at all probes with a constant frequency. We further focus on `traceroute` between RIPE Atlas probes and the 13 anycast addresses of root DNS servers [6].

We choose `traceroute` as it provides visibility into Starlink’s networks built with multiple PoPs, allowing us to derive round-trip time (RTT) to the Starlink *Carrier Grade NAT gateway* to the Internet (with address “100.64.0.1”). To derive the physical locations of such PoPs, we perform reverse DNS lookups of the probes’ public IP addresses. We collect public IP addresses by parsing the ‘src_addr’ field of the certificates obtained during RIPE “SSLCert” built-in measurements [23]. We choose this method instead of extracting them from the probe’s metadata because SSLCert experiments are recursive and run every 12 hours, allowing us to capture any temporal changes in the IP addresses. Reverse DNS lookups map a probe’s IP address with a hostname in the form of “customer.LOCATION.pop.starlinkisp.net”, thus revealing the location of Starlink PoPs. Each PoP represents a connection point between Starlink and the Internet backbone, serving potentially multiple “ground stations” or where LEO satellites connect to.

While analyzing RIPE Atlas dataset, we discovered that some probes were previously hosted by a different ISP than Starlink, and that the probes table only reports the most recent ASN. Further, a few probes are multi-homed and use mobile connections as a “failover”. To address these issues, we leverage `traceroute` data to verify that the (private) address of a Starlink Carrier Grade NAT gateway was present on the routing path.

3.2 Identifying Satellite Network Operators

The public datasets we used (M-Lab and RIPE Atlas) do not provide reliable information about the characteristics of access links their users or probes are connected to. Similarly, crowdsourcing platforms (e.g., Prolific) either do not provide such information or, when they do, it is not always reliable. Intuitively, SNO users and measurements can be identified based on the originating Autonomous System Number (ASN) available either in the measurement datasets or via a tester’s IP address. However, previous works have shown that accurate ASN-to-operator mapping information is lacking [68, 79]. For instance, an SNO might route the wireline traffic of its corporate offices via the same ASN used for its satellite customers. In addition, SNOs might rely on different technologies (LEO, GEO, MEO, as well as a mix) for which no public ASN-to-technology information is currently available. To tackle this challenge, we devise the following methodology (shown in Figure 1) which we

describe below.

ASN-to-SNO mapping. (steps 1a and 1b in Figure 1) We map a client’s ASN to its corresponding operator using ASdb [79] and Hurricane Electric’s (HE) BGP toolkit [4]. ASdb maps each ASN to its organization and a corresponding category (e.g., satellite or fiber) using machine learning approaches driven by manual data curation from various data sources such as Dun & Bradstreet and web classifiers like Zvelo [80]. Out of all 105k ASes recorded in ASdb, there are 129 ASes which are associated with the category “Satellite Communication” under “Computer and Information Technology”. However, we observe that several well-known SNOs like Starlink and Viasat are missing from the ASdb dataset. We fill this gap by searching for popular SNOs in the HE’s BGP toolkit which reveals an additional 35 ASes. In total, we obtain 164 ASN-to-SNO mappings.

Next, we gather information from IPInfo [44] about these 164 ASNs: organization, IP address ranges, and website. We then visit the website of each ASN present in our datasets (see Section 3.1) to augment the ASN-to-SNO mapping with access technology information. In the process, we discover that more than half of the 164 ASes do not necessarily belong to SNOs, instead to other similar providers such as Cable TV operators (e.g., Cable Axion), residential broadband (e.g., Filer Mutual Telephone), navigation services (e.g., Teletrac), teleport operators (e.g., United Teleports Inc), etc. After filtering them out manually, we obtain 67 ASNs belonging to 41 SNOs. (step 2 in Figure 1). Interested readers can refer to Table 3 in the Appendix for more information about these SNOs, their associated ASNs.

ASN-to-SNO validation. (step 3a in Figure 1) For the above 41 SNOs, we extract 21.3 million NDT speed test records from M-Lab traces collected between January, 2021 and April 2023 (see Section 3.1). For every NDT speed test, we take the 5th percentile of the latency estimated via TCP as an indication of the *access latency*. We then plot, for each SNO and ASN, the Kernel Density Estimation (KDE) [30] — a statistical technique used to estimate the probability density function of a random variable based on a set of observed data points — of its access latency.

Figure 2 shows representative KDE curves for the ASNs associated with Starlink and OneWeb (LEO), O3b (MEO), SES (both MEO and GEO), and TeleAlaska (GEO). The figure shows that LEO operators are characterized by low latencies, although with a significant difference between Starlink (median of 56 ms)




Figure 1: Data collection and processing pipeline: Methodology to identify Satellite Network Operators (SNOs) from the M-Lab dataset.




Figure 2: Kernel Density Estimation (KDE) curves by SNO showing standalone ASN for satellite user connections, Hybrid Access (H), Mixed Access (M).

and OneWeb (median of 154 ms). Note how comparable MEO latencies (280 ms for O3b and 220 ms for one component of SES) are to those of OneWeb, despite the latter being a LEO operator. GEO latencies, on the other hand, clearly depart from latency distributions of the other technologies, about 700 ms for both SES and TelAlaska. This analysis also unveils clear outliers at the AS level, such as ASN 27277 for Starlink whose KDE curve deviates from expected latencies of a LEO operator. Digging deeper, we find in the PeeringDB notes field of AS14593 (the AS used for Starlink customers) that AS27277 is used for Starlink’s corporate network [60], *i.e.*, its users are on an entirely terrestrial path allowing for lower latencies. Similarly, the KDE curve for ASN 201554 (SES) departs from its expected hybrid (MEO plus GEO) access – a bimodal distribution following a combination of the respective MEO/-GEO KDE curves. The same behavior is also observed within the same ASN, as shown by the low latency peak for the GEO-based TelAlaska [15].

IP prefix filtering. (step 3b in Figure 1) The previous analysis indicates that ASN filtering is not enough to identify network measurements associated with SNOs. Out of 41 SNOs, only measurements for OneWeb and O3b can be identified using only their ASNs. While some ASNs can be filtered as clearly departing from the expected latency pattern (e.g., 27277 for Starlink

and 201554 for SES), other ASNs are harder to filter as they exhibit mixed latency distributions within the same ASN (e.g., 10538 from TelAlaska). As a next step, we dissect the latency characteristics when considering IPv4 prefixes. In our dataset, there are two main IP blocks as per M-Lab annotation [52]: /24 (66.7%) and /21 (21.9%). Given it is the smallest and most popular IP block, we consider /24 IPv4 prefixes to group latency values. Note that the dataset only contains IPv6 data for four operators: HughesNet, SSI, TelAlaska, and GlobalSat. With the exception of HughesNet, this data collectively represents less than 0.1% of the full dataset. In case of HughesNet, we do not find any discernible pattern in terms of prefixes and latency. Therefore, we exclude IPv6 data from this filtering step.

Since no LEO-based SNO is left, we introduce only MEO (latency ≥ 200 ms, or 10th percentile from O3b latency distribution) and GEO (latency ≥ 500 ms [61]) filters. With this strict filtering, we retain /24 prefixes whose speed tests *only* exhibit latencies within the above filters, and which have at least 10 speed tests. This prefix-based filtering retains less than 1% of the remaining speed tests, and it spans 25 /24 prefixes from 6 SNOs (Figure 3a). The figure shows five GEO operators and one operator (SES) supporting both MEO and GEO. Note that SES acquired O3b, the only MEO operator in our dataset, in 2016 [70]; in the following,

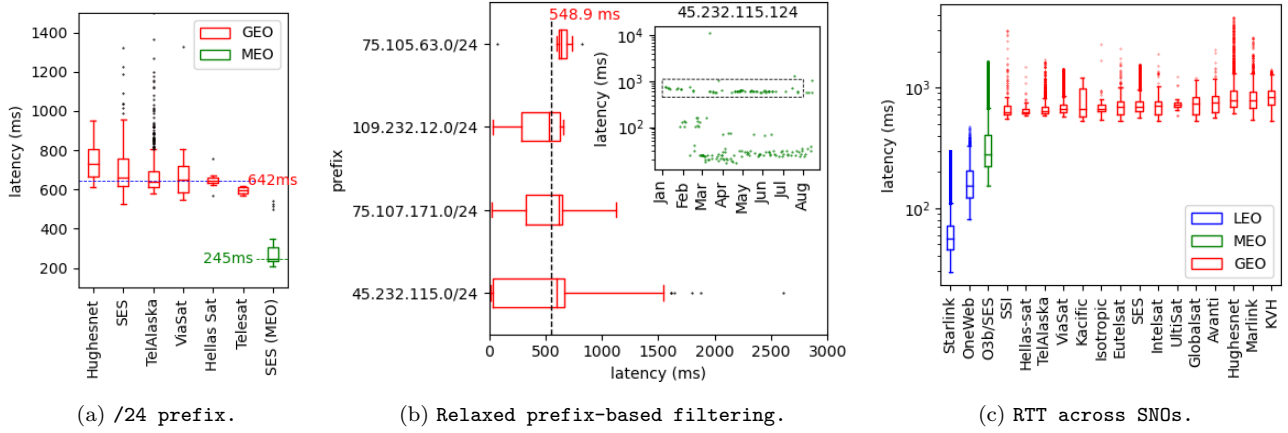


Figure 3: Analysis of the methodology we devised to identify SNOs.

we refer to O3b/SES as a combined MEO operator.

Relaxing the prefix-based filtering. (step 3c in Figure 1) Figure 3b shows the latency distribution for the remaining IP ranges after prefix filtering for Viasat. The figure shows that 75.105.63.0/24 was discarded just due to few outliers – thus indicating the prefix filtering might be indeed too strict – while the other three prefixes have GEO-like median latencies (around 500/600 ms) but also very wide latency distributions, e.g., 45.232.115.0/24 has 30% of latency values smaller than 70.6ms. The inset of Figure 3b shows the evolution over time of the latency values for one specific IP/user in the prefix 45.232.115.0/24. The figure shows three *clusters* of latency values centered around 600ms (GEO-like), 100-150ms, and 20-40ms. Clearly, such low latencies are incompatible with a GEO-based connection, and likely representative of a terrestrial connection for which the satellite link acts as a *backup* for some unreliable wired access offered by the same operator [76, 42].

Motivated by the above observation, we relax the prefix-based filtering, thus tolerating mixed latency patterns within /24 prefixes, or even for a single IP address. For the 6 SNOs covered by the prefix-based filtering (see Figure 3a), we further allow speed tests with latencies bigger than the minimum latency observed, e.g., 548.9ms for Viasat (blue-dotted line in Figure 3b). For the other SNOs, we use the minimum latency across the 6 SNOs covered by the prefix-based filtering (527ms).

SNO discovery in public traces. (final data accu-

mulation step 1 in Figure 1) After applying the above strategy on M-Lab traces (see Section 3.1), we identify 18 SNOs (2 LEO, 1 MEO, and 15 GEO) with a minimum of 34 speed tests (Kacific) and a maximum of 11.7 million measurements (Starlink), as summarized in Table 1. Instead, for RIPE Atlas, we identify 67 probes on Starlink, 6 on Viasat, and 2 on HughesNet. However, Viasat and HughesNet probes were all inactive during the past year or more according to RIPE Atlas website [11]. Hence, we restrict the RIPE Atlas analysis to Starlink only, whose probes are distributed across 15 countries. Table 2 shows the number of `traceroute` measurements per country, with the US having the highest number of measurements (about 3M), followed by Australia (460k). We add that due to the nature of the available data and the diversity of testers enrolled, we could not study all SNOs with equal depth. However, we have made every effort to conduct a thorough analysis with the available data.

3.3 Prolific Census

Prolific [10] is an online crowdsourcing platform with more than 130,000 vetted testers. Prolific offers powerful prescreening APIs for recruiting testers based on demographics, location, and even Internet Service Provider (ISP). ISP-based prescreening is only available for testers located in the US and the UK, where Prolific claims to be capable of correctly verifying ISP information [65].

According to Prolific’s prescreening at the time of this study, 160 testers subscribe to any one of three

SNO	# access	SNO	# access	SNO	# access	SNO	# access	SNO	# access
Starlink	11.7M	TelAlaska	3.05K	KVH	951	Avanti	122	Isotropic	35
O3b/SES	78.1K	OneWeb	2.95K	SSI	260	IntelSat	91	Kacific	34
Viasat	50K	HughesNet	2.80K	Eutelsat	235	Hellas-Sat	48		
SES	23.2K	Marlink	1.42K	GlobalSat	135	Ultisat	37		

Table 1: Filtered SNOs and the total number of access from each SNO. Starlink & OneWeb operates in the LEO orbit, O3b/SES in MEO, and the rest operate in the GEO orbit.

Country	# Probes	Start time	# trace-routes		# Probes	Start time	# trace-routes		# Probes	Start time	# trace-routes
AT	2	22/05	0.24M	DE	5	22/05	0.71M	NL	3	22/05	0.38M
AU	4	22/05	0.46M	ES	2	22/06	0.10M	NZ	1	22/05	0.22M
BE	1	23/01	0.07M	FR	5	22/11	0.35M	PH	1	23/03	0.02M
CA	2	22/05	0.28M	GB	5	22/08	0.29M	PL	1	23/01	0.06M
CL	1	23/02	0.05M	IT	1	22/10	0.12M	US	33	22/05	3.08M

Table 2: Summary of RIPE Atlas dataset.

SNOs: Starlink, HughesNet, and Viasat, *i.e.*, no other SNO is detected in their list of available operators. We recruited these 160 testers for a survey hosted on our server, where we ask them to provide the name of their provider, their location (city/state), and a score of how satisfied they are with their service. We ran the survey for seven days, attracting 30 participants. We found that only 20 testers were connecting from an IP address belonging to the above SNOs. Intuitively, it can be hard for Prolific to constantly verify how their testers are connected. For example, a tester can create an account at home but then participate in surveys from her phone or from work. This suggests a limitation in Prolific’s prescreening, but it also means that potentially more Prolific testers might have access to a satellite-based connection.

Based on the latter observation, we created multiple measurement campaigns over the course of two weeks, where we do not leverage Prolific prescreening, but request, in the study description, that only testers connected via an SNO participate. At the same time, we enable IP address-based access control at our server (via nginx [66]) which allows only IP addresses belonging to SNOs from Table 1. Our studies attracted a total of 14,371 Prolific testers, out of which 57 were actually connected via Starlink, HughesNet, and Viasat. We find that, overall, Starlink users are much more satisfied with their service than both HughesNet and Viasat users (see Figure 14 in the Appendix). For example, only one user out of 20 considers Starlink connectivity as “poor”, while most users consider it either good or very good. On the other hand, “ok” is the highest score reported for both HughesNet (55% of the answers) and

Viasat (18% of the answers).

3.4 Limitations

Lack of ground truth. The accuracy of the methodology described in Section 3.2 to identify measurements performed via SNO is hard to quantify. This is because we lack ground truth to compare against, and we thus rely on a comparative analysis between operators latency profiles derived from M-Lab data. This means that while our methodology catches obviously incorrect associations between ASNs and SNOs (see Figure 2a for example), as well as IP reuse across non-satellite access (see Figure 3b), it can potentially introduce some errors especially when the latency difference between technologies is small such as when comparing LEO and MEO. However, our analysis in Section 3.2 shows that LEO operators are easy to identify, and most of the remaining operators are GEO-based, whose latency profile is hard to confuse with LEO or terrestrial accesses due to the long latency inherent to GEO-based connectivity. We thus believe that our methodology catches most major errors, but we acknowledge that further unknown errors might exist. We opt not to perform any geographical analysis of end users using the M-Lab data due to the known fallacies of geo-location mapping [62, 39, 27]. Further, this introduces an additional layer of complexity in our data handling process. However, we perform spatial dissection for data collected from RIPE probes and Prolific users, as their locations are known.

Speedtest accuracy. NDT measurements on M-Lab rely on a single TCP connection, and it is well-

known [35] that single-flow TCP often under-estimates the available bandwidth due to slow-start. Further, owing to frequent satellite handoffs, it is possible for LEO satellite communication networks to have temporary periods of high packet loss which are not related to congestion [48]. If such losses occurred during the NDT tests in our dataset, it would have led to the TCP flow decreasing its sending rate and consequently underestimating the bandwidth. This is a known artefact of TCP-based speed-tests, but from our dataset it is not possible to make an ex post facto analysis to filter out tests which suffered losses due to satellite handoffs etc. In addition we focus exclusively on download tests for two reasons. Firstly, the kernel level `TCP_Info` is solely captured at the server, ensuring the reliability of the values for download measurements. Secondly, the upload and download traces are obtained separately, necessitating the use of a time-interval based heuristic to merge the upload and download tests from the same client [59, 74]. As also noted in [45], we acknowledge that users typically run speed tests when they experience degraded network conditions. It follows that results derived from M-Lab speed tests might underestimate the actual network conditions. We want to highlight that the presence of middleboxes, VPNs, NAT, or any other end-user configurations (such as speed throttling at the browser) could potentially affect the results we obtain. However, we argue that such variations are unavoidable in real-world data traces, and these data anomalies should be considered as outliers that may not necessarily have a significant impact on the overall trends.

4 A Bird’s-Eye Look Into SNOs

Latency distributions. As shown in Figure 2, SNOs’ performance is expected to vary depending on the orbit they operate in. Figure 3c shows boxplots of the (access) latency per SNO using the M-Lab dataset. As expected, the latencies for GEO SNOs are higher, with a median of 673.5 ms, followed by MEO (279.36 ms) and LEO (56–154ms). Among the two LEO operators, Starlink is overall 3x faster than OneWeb. When focusing on GEO operators, the figure shows a significant difference between the best performing GEO SNO (SSI, with a median latency of 620.4 ms) versus the worst performing one (KVH, with a median latency of 835.2 ms).

Next, we analyze the evolution over time of the access latency for the 5 most popular SNOs per orbit: Starlink, OneWeb, O3B, HughesNet, and Viasat. Fig-

ure 4a shows the median access latency per day and operator over the last year. The figure confirms the trend observed in Figure 3c, at any point in time. It further shows that Starlink and Viasat, *i.e.*, the two most prominent operators in their respective orbits, are quite stable, with daily latency variation (95th %ile) up to 3.1% and 7.2%, respectively. O3b shows a similar trend, although its daily variation increases to 41.4%. Both HughesNet and OneWeb are instead often affected by significant daily latency variations, up to 72% for HughesNet and 120% for OneWeb.

Variability in latency. Next, we investigate the impact of variability in latency (*jitter*) when accessing SNOs operating in different orbits. In order to take into account the large differences in the overall latencies between LEO, MEO and GEO, we normalize the jitter (95th percentile) relative to the 5th percentile latency for each M-Lab request, showing how the jitter varies as a fraction of the lowest latencies achievable in each orbit. Figure 4b shows Cumulative Distribution Functions (CDFs) of this *jitter variation* per orbit. The figure shows that LEO exhibits a greater degree of jitter variability compared to GEO, with a median of 0.5 as opposed to 0.28. This observation indicates that, even though GEO SNOs are characterized by overall much higher latencies than LEO SNOs, they are overall more stable. This discrepancy can be attributed to the fact that LEO satellites require frequent satellite handoffs [58, 22]. MEO is characterized by less frequent satellite handoffs than LEO [58] and shows jitter variability similar to GEO. However, the figure also shows that, for 10% of the values, MEO’s jitter variation is comparable to and even higher than LEO. This suggests that, when there is a handoff, recovering from it is more difficult in MEO networks; this happens because there are fewer satellites deployed than in the LEO satellite mega constellations. Note that, while this demonstrates the network’s variability in connection based on latency, we observe that LEO outperforms GEO when comparing absolute jitter values, as shown in Figure 4b (inset) — over 80% of the GEO trace exhibited a jitter of 100ms or more, while LEO had less than 20% with such high jitter.

Retransmissions. Finally, we shift our attention to a major side-effect of high latency and variability in latency: packet retransmissions. Figure 4c shows the percentage of bytes that required retransmission. LEO satellites have fewer retransmissions than MEO, which is expected considering that the lower latencies to LEO

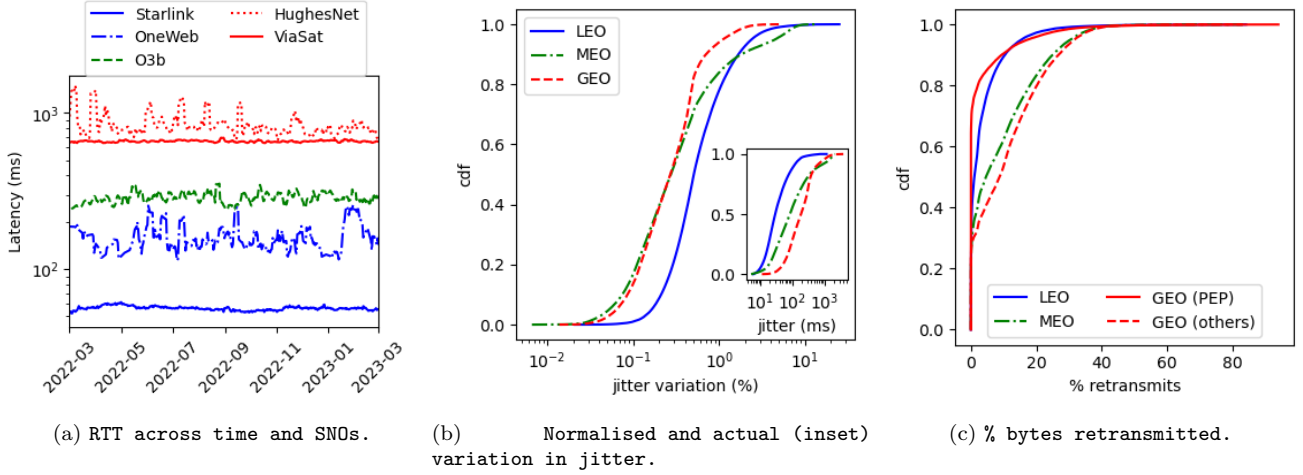


Figure 4: Performance of LEO, MEO and GEO SNOs.

orbits can lead to fewer retransmission timeouts and errors. The picture for GEO operators is more nuanced. One class of operators, shown as “GEO (others)” have a significant (median 8.74%) fraction of data retransmitted, which can be expected in GEO, given the difficulties of transport protocols on high latency asymmetric links. However, four GEO operators (HughesNet, Viasat, Eutelsat and Avanti) are exceptions to this trend and their retransmission fractions are close to those of LEO operators. Their retransmission factor CDF is shown as the line “GEO (PEP)” in Figure 4c, as we notice that all four operators use Performance Enhancing Proxies (PEPs) [69, 72]¹. PEPs manage the TCP connection on either side of the high latency bent-pipe link and thereby mitigate the effect of high latencies leading to retransmission timeouts. If there are any latency-related retransmissions on the satellite link, these can be managed by the PEPs so that they are transparent to the TCP connection between the end hosts being connected [40]. Figure 4c points to the effectiveness of this strategy.

Geographic connectivity characterization. End-to-end latency in an SNO can be influenced not just by the orbit it operates in, but also how well its network infrastructure on the ground is provisioned, e.g., the number/location of its PoPs and their physical connectivity to the SNO’s upstream provider networks. However, mapping the geographic coverage of an SNO can be challenging. For the majority of SNOs, there is no publicly available infrastructure map. As of this writing, we only

find PoP location information for Starlink [17], SES [13] and Hellas-Sat [1]. In addition, from reverse DNS lookup of 900K SNO IP addresses discovered from M-Lab datasets, we find that no SNO other than Starlink encodes PoP location information in subscriber domain names. Finally, the client location information available in M-Lab datasets may not necessarily represent the PoP locations of SNOs, but simply the (approximate) locations of clients inferred by Google’s location service.

Given these challenges, our approach is to “indirectly” infer the geographic coverage of SNOs from BGP peering data. Our intuition is the following. Since none of the existing SNOs is a tier-1 ISP, they must peer with larger upstream ISPs to obtain global reachability of their ground infrastructure. If an SNO peers with upstream ISPs in many different locations, that should be an indication of similar geographic spread of its ground stations. Based on this intuition, we analyze BGP peerings of SNOs with the BGP route-views [16] collected from similar periods as our SNO datasets.

Initially, we tried to identify exact physical locations of SNO’s BGP peering points at available Internet exchanges [9]. However, we find that SNO information at public exchange points is not very complete except for very few SNOs with many BGP peers, such as Starlink and SES. For broader coverage, we end up using peering ASes’ country jurisdiction as the approximate peering location. AS country information is from the Regional Internet Registries, where the ISO 3166 country codes are added during ASN assignment [43]. When we cross-check the inferred country-level PoP locations

¹HughesNet [2], Viasat [77], Eutelsat [49] and Avanti [33].

with available ground-truth PoP location data of several SNOs, we find that our approach discovers 10 out of 30, 7 out of 22, and 2 out of 2 countries for Starlink, SES and Hellas-Sat, respectively. The discovered countries cover 74%, 57% and 100% of their city-level PoP locations, respectively. In case of Starlink, missing country-level PoP locations include European countries (e.g., Spain, Portugal, France, Norway), South American countries (e.g., Mexico, Brazil, Chile, Peru) and Caribbean countries (e.g., Fiji, Dominican Republic). Upon closer inspection of the discrepancies, we find that some of Starlink’s peering ASes actually have *continent-wide* presence in their network connectivity (e.g., AS1299 (Arelion) and AS6762 (Telecom Italia Sparkle) for Europe and AS7195 (EdgeUno) for South America). Hence, Starlink may peer with these neighbors in multiple country locations within the continents, beyond the single jurisdiction registered with these AS numbers. In fact, multi-location peerings are reported in CAIDA’s peering geolocation data set [50], where we observe that about 6% of peering links are associated with more than one location. Unfortunately, we could not use the CAIDA dataset because its collection period (2016) is considerably earlier than ours (2021–2023), and quite a few new ASes have appeared since then. Although our approach can underestimate the geographic coverage of SNOs due to this limitation, its main benefit is that it enables comparative analyses of geographic coverage of *all existing SNOs*, as well as longitudinal studies on the *historical evolution* of a specific SNO’s ground infrastructure.

Figure 5 visualizes BGP peerings of few SNOs; more examples can be found in Figure 12 in the Appendix. For a given SNO, the label of each peering AS indicates its ASN and country. The diameter of each peering AS indicates its “size”, estimated with its node degree. From relative AS sizes, we can speculate whether or not a peering neighbor is an SNO’s upstream provider [37]. For example, Starlink is connected to AS3356 (Level3) which is much bigger than Starlink, so this AS is likely a Starlink’s upstream provider. Conversely, Kacific peers with much smaller ASes than itself, which we confirm are small regional ISPs obtaining satellite access from Kacific [5].

From geographic characterization of SNO PoPs, we make the following observations. First of all, the upstream connectivity of SNOs in both LEO and GEO categories is significantly varied. For example, in GEO category, Hellas-Sat and UltiSat are not connected to any tier-1 providers, while Viasat and Kacific are well-connected to multiple tier-1 providers. Similarly, in

the LEO category, Figure 5 shows that Starlink is peering with major upstream providers all around the globe, while OneWeb is connected to only two US-based providers. This could potentially explain the significant latency difference between them in Figure 3c. We also observe (not shown) that Starlink shows consistent latency performance worldwide, while OneWeb exhibits skewed performance in North America vs. the rest of the world. The SNO in MEO category (e.g., SES) tends to be much better connected to major transit providers than pure GEO-based SNOs. MEO-based SNOs are likely to pursue more aggressive peering to cover more grounds as they expand to LEO satellites. Finally, we observe that different SNOs exhibit widely varying growth trajectories in their PoP infrastructures over time, e.g., Starlink had explosive growth compared to HughesNet which was stagnant between 2021 and 2023 (Figure 13); such geographic diversity may help Starlink as it expands to new countries [71]. We explore Starlink latencies in different countries next.

5 A Closer Look at Starlink

This section zooms in on Starlink using built-in measurements performed by 67 Ripe Atlas probes between May 2022 and May 2023 (see Table 2). Our goal is to shed some light on Starlink infrastructure worldwide and its impact on network latency.

Rest of the world.² Figure 6a shows boxplots of the RTT measured over one year between RIPE Atlas probes located outside the US (34 probes) and Starlink PoPs; see Table 2 for the number of measurements per country. Results are organized by country and grouped by continent using the following color scheme: green for Europe, blue for North America (Canada only), yellow for South America, purple for Asia, and red for Oceania. The figure shows that New Zealand and Chile probes exhibit the lowest RTT to their respective Starlink PoPs, with median values at approximately 33 ms. Most European countries follow closely with median values of 35–40 ms. Canada and Australia experience slightly higher RTTs, averaging at around 45 ms. The highest RTT to a Starlink ground station is in the Philippines, with a median of 80 ms, almost twice what measured in other locations.

To investigate the cause behind this substantial RTT disparity, we analyze the geographical placement of

² “Rest of the world” refers to worldwide, excluding the United States.

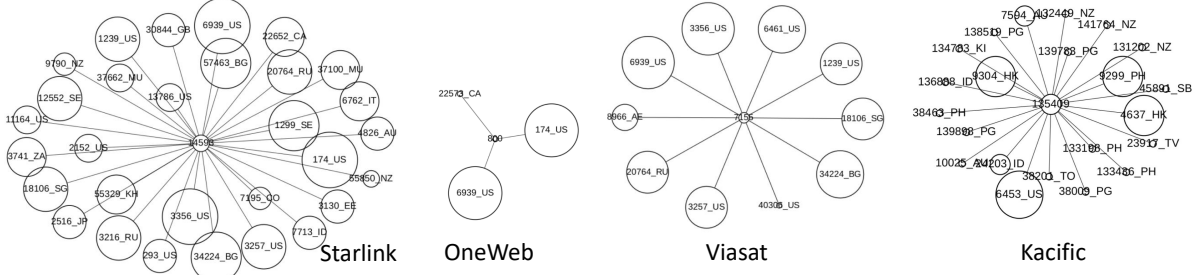


Figure 5: BGP peering visualization of SNOs (based on BGP route-views collected on 2023/1/1).




Figure 6: Analysis of `traceroute` measurements between Ripe Atlas probes on Starlink and root DNS servers; Rest of the world (34 probes, in total).

Starlink PoPs used by each probe. Figure 7 visualizes the location of RIPE Atlas probes (yellow circles) along with the PoPs (black circles) they connect to distinguishing between the US, Europe, Oceania, and Philippines. The green lines show active connections between a probe and a PoP, whereas, the red lines show inactive (*i.e.*, previously used but not at the time of writing) connection between a probe and a PoP. We do not show the probe-PoP pair in Chile due to space limitation and since they are co-located (approximately 75 km away), thus explaining the very low latency. When focusing on the Philippines probe, located in Manila, the figure shows that the PoP for this probe is located in Tokyo, Japan (“customer.tkyojpn1.pop.starlinkisp.net”) likely causing the extra delay (roughly 40 ms). While we cannot measure the link between Starlink antennas – located in the Philippines according to Google Maps – and the Tokyo PoP, we checked the RTT from Manila to Tokyo [78] and found fairly high values, 177 ms on average.

We now move to the the analysis of RTT to root DNS servers (see Figure 6b). The figure shows negligible difference between European countries, with median RTT values comprised between 40 and 49 ms (apart from

Spain that has a slightly higher median RTT of 58 ms); this is expected given their comparable RTT to their respective Starlink PoPs and the presence of many root DNS servers in Europe. Despite having the lowest RTT to a Starlink PoP, the Chilean node experiences an extra 10-20 ms delay to root DNS servers making its performance comparable to the Canadian probe. This happens because only 7 out of 13 root DNS servers are present in Chile, requiring longest routes for about 50% of the requests, which also explain the wide distribution of RTT values. This behavior is verified by Figure 6c, which shows a wide variability of path to root DNS servers comprised between a minimum of 5 hops (*e.g.*, to the L root accessed in Santiago, Chile) and 20 hops or above (*e.g.*, to the M root which is currently not present in South America). A similar behavior applies to New Zealand and Australia, which only host few local root DNS servers and, for most queries, requires 100-150 ms and more than 10 routing hops. Philippines trails with the highest delay to root DNS servers (about 200 ms).

United States. Figure 8a shows boxplots of the RTT measured over one year between RIPE Atlas probes located in the US (33 probes) and Starlink PoPs. We



Figure 7: Location of RIPE Atlas probes and connectivity to their Starlink PoPs. Green line refers to active links as of April 2023. Red dotted lines refer to previous links observed between April 2022 and April 2023.

only focus on this metric since it is the most relevant to this analysis. Results are organized by state and grouped as follows: Northeast, Southeast, Central, East North Central, South, Southwest, West, Northwest, and Alaska. If we focus on the median, the figure shows that the minimum RTT to Starlink PoPs sits around 45 ms for the following states: Oregon, Washington, Virginia, New York, and Pennsylvania. Note that this is about 10 ms higher than what observed for the best performing countries in the rest of the world (see Figure 6a). The remaining states, with the exception of Alaska, show slightly higher median RTT values, with Arizona reaching a maximum of 55 ms.

Alaska exhibits a significantly higher RTT than all the other states, with median RTT of 80 ms (75th percentile RTT of 120ms). To understand the underlying cause of this high RTT, we again resort to the analysis of the geographical location of the probe-to-PoP mapping. Figure 7a shows that the Alaska probe is connected to a PoP in Seattle (Washington State), approximately 2,697 km away, which is likely responsible of the additional RTT incurred by the probe.

Finally, we analyze Starlink latencies over *time*. For most probes, we observe non statistically significant variations of RTT to their PoPs over one year. Figure 8b visualizes six probes which instead see a significant change in their Starlink latencies over time. The figure shows that the New Zealand probe experienced a 20 ms latency reduction since July 12th, 2022 when its PoP changed from Sydney (Australia) to Auckland (New Zealand), as indicated by the red dotted line in Figure 7c. Indeed, such PoP change from Sydney to Auckland was tested between June 2nd and 3rd (2022), which is reflected in the sudden latency reduction around the beginning of June. A similar obser-

vation can be made for the Netherlands probe, where a 10 ms latency drop was due to a shift from a PoP in Frankfurt to one in London, as shown in Figure 7b. Finally, the figure shows also the potential “damage” associated with a PoP change. In the case of Nevada, in September 2022 the PoP was changed from Los Angeles to Denver (see Figure 7a), which caused a 2x increase of the RTT to the assigned PoP. The change was reverted about one month later, allowing the RTT to return to its previous values for one of the two probes in Nevada. No PoP change was detected for the other probe. However, the figure shows a high variable, although down trending, latency for this probe which eventually stabilize around 55ms.

6 Applications Performance

Finally, we expand on how *real users* perceive the performance they obtain from SNOs by conducting a user study on real customers of Starlink, Viasat and HughesNet.

6.1 Methodology and Data Collection

Browser add-on design. As in [48, 75], we rely on a Chromium add-on to perform Internet measurements from end users in the wild. This design choice is motivated by its ease of installation and wide compatibility, as Chromium-based browsers account for 80% of the market [73]. The add-on guides the tester through a set of experiments using local “hooks” implemented in the add-on, providing instructions and requesting the launch of the next measurement as the tester is ready. After each experiment, the add-on notifies the tester to close the tab associated with an experiment, if any, and re-

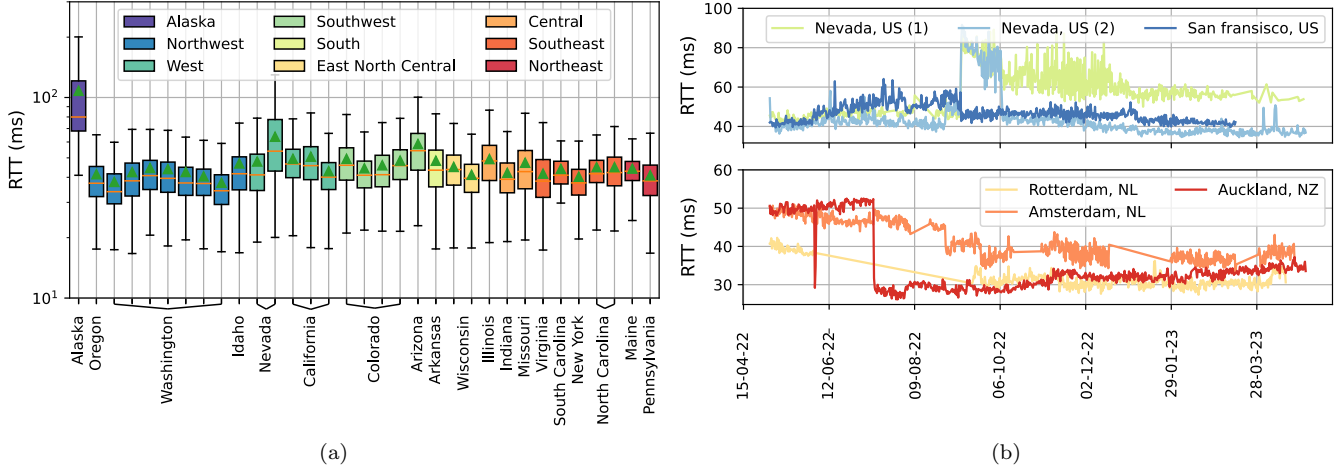


Figure 8: (a) RTT between the 33 RIPE Atlas probes located in the US and Starlink PoPs; and (b) RTT over time between (New Zealand, Netherlands, San Francisco, and Nevada) RIPE Atlas probes and Starlink PoPs.

turn to the addon. The addon also communicates with our server (the “experiment manager”) to be instructed on what experiments should be run, and to report the data collected. Upon completion of a study, the tester is presented with a code which is synchronized (using Prolific’s API [63]) with the respective job posted on Prolific.

We have developed code hooks in the addon for four measurements: Content Delivery Network (CDN) measurements, speedtest (using *fast.com*), webpage loads, and video streaming (using YouTube). The speedtest experiment relies on the code provided by [75], but with slight adjustments to enable the upload bandwidth measurement at *fast.com*. When needed, JavaScript (JS) code is injected into a new tab to collect relevant statistics, e.g., performance metrics when loading a webpage, and to help a tester, e.g., inform when a measurement is completed. This injected JS code also communicates with a background script which has permission to: 1) communicate with the experiment manager, 2) collect screenshots (when needed).

After installation, the addon informs the experiment manager of a new tester. At this time, the tester is asked to wait for one minute during a *warm-up* while the background script measures the “clock drift” of the client’s local clock (along with RTT to our server), and identifies the installed DNS resolver using <https://test.nextdns.io/>. Next, the addon measures the performance of several CDNs (Cloudflare, Google, jsDelivr, StackPath, and Fastly) by fetching a popular JS file (*jquery.js*) while measuring download time, and

collecting HTTP headers. Note that jsDelivr advertises “optimal” speed by matching each request to an optimal CDN based on uptime and performance [47]. Finally, the same CDN measurement – but fetching the minified version of this script (*jquery.min.js*) – is repeated meanwhile the tester is asked to provide his/her current location (city and state).

To avoid interference with the browser cache, we request older versions of each file (regular and minified), ranging from 3.0.0 up to 3.6.4, verifying that the file is not served from the cache; if it is, we attempt a different version in this range until no cache entry is found. This allows to perform multiple experiments with the same tester, and to handle the unlucky event of this file being previously stored in the cache. We further discard the first file download to realize a DNS primer. The maximum file size variation observed across versions is about 3 KB. When comparing CDN providers, Cloudflare serves the most compressed versions of both files, *i.e.*, 28 KB (versus 31-33 KB) for the minified version, and 71 KB (versus 86-89 KB) for the regular version.

After the initial warm-up, when the tester is ready they can click on a button in the addon which will open a tab pointing to *fast.com*. We use the injected JS code to monitor the statistics on screen, such as bandwidth, latency, and client/server location, while also detecting when the measurement is completed. The next experiment involves loading multiple webpages sequentially, based on a list provided by the experiment manager. The addon automatically iterates through each URL waiting for the *onload* event (or a timeout of about

60 seconds) while collecting and reporting performance metrics. As shown in Figure 10b, this timeout was only triggered by one tester from HughesNet (PLT of 62.6 seconds). One single screenshot is taken before leaving the tab, which helps verify the data collected.

In our data collection, we instrument the addon to load two demo pages provided by Akamai to compare the performance of HTTP/1.1 (H1) [20] and HTTP/2 (H2) [21]. Our rationale is threefold. First, these pages have constant sizes, no impact of personalized ads, and they bypass the browser cache using unique Web object identifiers. Second, these pages are hosted by a large CDN with a wide footprint and thus at a relatively short distance to all our users. Third, they are *designed* to allow a comparison of the performance of H1 and H2.

The webpage test is followed by video streaming, where a predefined YouTube video is played for 60 seconds. During video playback, injected JS code monitors the YouTube’s “Stats-for-Nerds” which reports statistics like video quality, data available in the buffer, and number of dropped frames. To avoid interference with the browser cache for returning testers, we use a 10 hours 4K video and increase the starting playback time by one hour, e.g., by appending `&t=3600s` to the URL.

Data collection. During April-May 2023, we launched Prolific measurement campaigns targeting the 56 testers we previously identified as SNO subscribers. Over one month, we recruited 20 testers (10 on Starlink, 5 on HughesNet, and 5 on Viasat) willing to install and run our addon once a week (*i.e.*, four runs in total) on random days and times. The remaining 36 Prolific testers did not accept our jobs. Some had temporarily no access to their SNO, others were not willing to install software on their machines, and the rest just never responded to our job requests and direct messages.

6.2 Results

Speedtest. Figure 9 summarizes the speedtest analysis per SNO (Starlink, HughesNet, and Viasat), and per continent (North America, Europe, and Oceania). We report on both download and upload speeds, and network latency measured between each tester and the speedtest service provided by *fast.com*. Overall, Starlink offers much higher speeds in both download (70-150 Mbps) and upload (6-21 Mbs). While Viasat and HughesNet achieve comparable upload speeds (3 Mbps), Viasat offers much higher download speeds (10-40 Mbs), in line with what is advertised [19], while HughesNet

testers never experience more than 3 Mbps, which is far from the download speed advertised by this provider (25 Mbps [19]). With respect to the per-continent analysis, Starlink testers reach similar download speeds (~ 80 Mbps) in both North America (nine testers) and Oceania (only one tester located in New Zealand). Much higher speeds (median of 150 Mbps) are instead reached by five European testers (Italy, UK, Netherlands, and Czech Republic). A similar trend is observed for upload speeds as well, with European testers reaching the fastest speeds (~ 21 Mbps), followed by New Zealand with a stable speed of 13 Mbs, and North America with more variable speeds and a much lower median upload speed of 6 Mbps.

We now focus on the RTT between Prolific testers and *fast.com* servers. Figure 9c shows that, despite Viasat and HughesNet both being GEO networks, Viasat manages to save about 100 ms, *i.e.*, median latency to *fast.com* of 600 ms compared to 720 ms with HughesNet. Starlink testers benefit from the low latency offered by LEO satellites, with median latencies of 35.0 ms (North America), 38.0 ms (Europe) and 49 ms (New Zealand). These latencies are comparable with the RTT measured between RIPE probes and Starlink PoPs (see Figure 6a and 8a), indicating that *fast.com* servers are likely co-located with these PoPs. Note that the high latency values measured in Europe (up to 150 ms) are all associated with runs from the London tester. Given the results from Figure 6a, this latency represents an outlier due to some issues on the test side, e.g., bad WiFi setup.

Web browsing. We now focus on measurements related to Web browsing. Specifically, we report the time taken to download the popular *jquery* library from multiple CDNs, the time to perform DNS lookups, and how quickly Akamai test pages are loaded using both HTTP/1.1 (H1) and HTTP/2 (H2). Figure 10a shows, per SNO, boxplots of download time of *jquery*’s minified version (*jquery.min.js*) via five CDNs: Cloudflare, Google, jsDelivr, StackPath, and Fastly. For all three SNOs, Fastly provides the fastest downloads, with a median of 127 ms (Starlink), 950 ms (HughesNet), and 1,036 ms (Viasat). When focusing on Starlink, jsDelivr is the second fastest CDN with a median download time of 170 ms; when we analyze the response headers, we find that all jsDelivr requests are handled by Fastly. This result suggests that jsDelivr is effectively identifying the best performing CDN; however, this process requires one extra RTT which eliminates the benefit




Figure 9: Analysis of speedtest data (*fast.com*) organized by continent (North America, Europe, Oceania) and satellite-based operator (Starlink, HughesNet, Viasat).

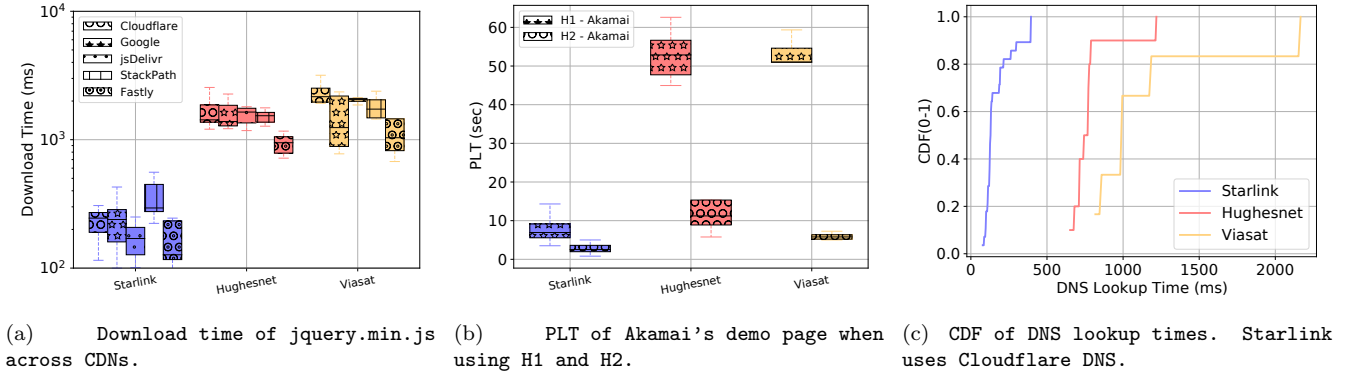


Figure 10: Web browsing analysis: Starlink vs. HughesNet vs. Viasat.

for GEO SNOs. For example, jsDeliver adds an extra 700 ms (download time of 1,641 ms) for HughesNet, making it slower than Cloudflare (1,427 ms), Google (1,385 ms), and StackPath (1,537 ms).

Although not shown to avoid cluttering the figure, a similar trend can be observed for `jquery.js`, *i.e.*, the non-minified version with a size of about 87 KB (versus 32 KB). In this case, Fastly serves the file in 190.0 ms (Starlink), 1,450 ms (Viasat), and 1,620 ms (HughesNet). Given the impact of JS libraries like jquery to start rendering webpages, this result showcases the importance of careful webpage development when dealing with testers behind such high latencies. A similar observation holds also for Figure 10b which reports the “page load time” of Akamai’s demo page when comparing H1 and H2. Adopting H2 is paramount for GEO testers, enabling performance comparable to using H1 on Starlink. Notice also the compound effect of the lower latency provided by Viasat (about 100 ms according to *fast.com*, see Figure 9c) when loading complex pages (hundreds of

small objects in the case of this demo) which can speed up webpage loads by multiple seconds.

Finally, Figure 10c shows the CDF of the DNS lookup time across SNOs. For this analysis, we leverage both the data collected by the addon when loading Akamai demo pages, and data collected by our server when loading the census form (see Section 3.3). We have verified, using <https://test.nextdns.io/> as discussed in Section 6.1, that all Prolific testers rely on what provided by their SNO, *i.e.*, Cloudflare for Starlink and their own DNS service for both both Viasat and HughesNet. Note that while DNS entries can be cached along the path, *i.e.*, browser, OS, router, and resolver, we minimize this issue by relying on unpopular domains (Akamai demo and our own domain) with short TTL (300 seconds). We further filter DNS lookup times smaller than the minimum RTT measured for each SNO. Figure 10c shows median DNS lookup times of 130 ms (Starlink), 755 ms (HughesNet), and 985 ms (Viasat). Given the overall faster RTT offered by Viasat, this result suggests

that HughesNet-provided DNS is faster than Viasat, and that its testers rely on default setting, *i.e.*, they did not manually set DNS in their devices. We would otherwise expect faster DNS lookup times for Viasat if a cloud-based DNS service (e.g., Google or Cloudflare) was used, since in this case the lower RTT offered by Viasat would dominate. We further verified this claim by messaging HughesNet and Viasat testers on Prolific.

Video streaming. Figure 11 summarizes the video streaming analysis per SNO reporting three metrics (download speed, buffer health, and percentage of frames dropped) as a function of the video quality experienced. We express the video quality as *megapixels* derived as “width” multiplied by “height” from the video resolution. For example, a 1080p video (1920x1080) corresponds to about two megapixels. This step is necessary since we observe 18 unique video resolutions, which would make the figure unreadable. Further, we report the median video quality for each experiment since the YouTube player is adaptive to network conditions. Note that the test video has a maximum quality of 2160p (3840x2160) or about 8 megapixels.

Overall, Figure 11 shows that only Starlink testers are able to play the video at high resolution (2 megapixels or higher). On the other hand, for both HughesNet and Viasat, about 0.5 megapixels (*i.e.*, less than the popular 360p, or 480x360) is the most common resolution experienced by their testers, although Viasat testers manage to maintain 60 frames-per-second (FPS) and sometime venture in higher resolutions of about 1 megapixel. This is due to a lack of download bandwidth as shown in Figure 11a (download speed measured by YouTube) which confirms the trend observed in Figure 9a. Indeed, 1080p or higher is hard to achieve also for Starlink testers, requiring sacrificing *buffer health*, *i.e.*, the amount of seconds of video available to be played.

Figure 11b shows how most runs, regardless of the SNO, are characterized by buffer health of between 40 and 65 seconds, whereas this number drops to 15 and 30 seconds when considering high resolutions for Starlink testers. Conversely, the percentage of frames lost is less correlated with the video quality, e.g., no frame was lost for Starlink runs where the video was played at full resolution (4K or more than 8 megapixels). These overall high losses rate are more likely due to the frequent handovers (for Starlink) as previously discussed in Section 4. Note that YouTube does not directly report the number of video stalls detected. However, this metric can be approximated by the likelihood of the buffer

health to reach zero. Out of 56 testers, only 4 testers (2 on Starlink and one on Viasat and HughesNet) have experienced some video stalls, between a minimum of 5% and up to 32% of the video playback.

7 Conclusion

By *opportunistically* utilizing a wide variety of publicly available data sources together with user studies, this paper performs what we believe is the first study which compares the performance of all three kinds of SNOs (LEO, MEO and GEO) and identifies factors critical for their performance. We observe that LEO networks experience greater relative variations in jitter, which can impact applications that rely on a consistent latency profile. Furthermore, we identify a number of potential factors that web applications can optimise, such as the choice of CDN used etc., as well as factors under SNO control, such as the location and geographic distribution of PoPs, which can each help improve global performance of satellite-based networks. We believe that these findings will be valuable to various stakeholders, from application developers looking to design their products with next-generation satellite access in mind, to CDN and satellite network operators seeking to make careful technology selection (for example, choosing the right locations to peer with SNOs). The study also highlights the benefits of bringing together different public data sources (M-Lab, RIPE, CAIDA) to gain a comprehensive understanding of connectivity performance in satellite networks.

As a part of future work, we plan to investigate temporal trends and assess the underlying causes of latency fluctuations over time among different operators. Additionally, we plan to conduct a more in-depth analysis of TCP traces to thoroughly examine retransmission rates and study the performance characteristics of satellite internet service providers.

Acknowledgments

This work is supported by EU Horizon Framework grant agreement 101060294 (XGain) and the ISOC Pulse Research Fellowship awarded to Aravindh Raman.




Figure 11: YouTube analysis: Starlink vs. HughesNet vs. Viasat.

References

- [1] Hellas-sat satellites & ground infrastructure. <https://www.hellas-sat.net/coverage>, 2023.
- [2] Hughes: Datasheet. <https://www.hughes.com/sites/hughes.com/files/2022-03/JUPITER-System-DVB-S2X.pdf>, 2023.
- [3] Hughes network systems. <https://www.hughes.com>, 2023.
- [4] Hurricane electric bgp toolkit. <https://bgp.he.net/>, 2023.
- [5] Kacific authorized distributors. <https://kacific.com/kad>, 2023.
- [6] List of root servers. <https://www.iana.org/domains/root/servers>, 2023.
- [7] M-lab bigquery quick start. <https://www.measurementlab.net/data/docs/bq/quickstart/>, 2023.
- [8] Measurement lab (m-lab). <https://www.measurementlab.net>, 2023.
- [9] Peeringdb. <https://www.peeringdb.com>, 2023.
- [10] Prolific. <https://www.prolific.co>, 2023.
- [11] Ripe atlas. <https://atlas.ripe.net>, 2023.
- [12] Ripe atlas docs — built-in measurement. <https://atlas.ripe.net/docs/built-in-measurements/>, 2023.
- [13] Ses network map. <https://www.ses.com/our-coverage/teleport-map>, 2023.
- [14] Starlink. <https://www.starlink.com>, 2023.
- [15] Telalaska teleport. TelAlaska, 2023.
- [16] University of oregon route views project. <https://www.routeviews.org/routeviews/>, 2023.
- [17] Unofficial starlink global gateways & pops. https://www.google.com/mymaps/viewer?mid=1805q6r1ePY4WZd8QM0aNe2BqAgFkYBY&hl=en_US, 2023.
- [18] Viasat. <https://www.viasat.com>, 2023.
- [19] Viasat vs hughesnet: Which satellite internet provider is better? <https://www.satelliteinternet.com/resources/viasat-vs-hughesnet>, 2023.
- [20] Akamai. Http/1.1 demo page. http1.akamai.com/demo/h1_demo_frame.html, 2023.
- [21] Akamai. Http/2 demo page. http2.akamai.com/demo/h1_demo_frame.html, 2023.
- [22] Ian F Akyildiz, Hüseyin Uzunalioglu, and Michael D Bender. Handover management in low earth orbit (leo) satellite networks. *Mobile Networks and Applications*, 4:301–310, 1999.
- [23] Vaibhav Bajpai, Steffie Jacob Eravuchira, and Jürgen Schönwälder. Lessons learned from using the ripe atlas platform for measurement research. *SIGCOMM Comput. Commun. Rev.*, 45(3):35–42, jul 2015.
- [24] Vaibhav Bajpai, Steffie Jacob Eravuchira, and Jürgen Schönwälder. Dissecting last-mile latency characteristics. *ACM SIGCOMM Computer Communication Review*, 47(5), October 2017.

- [25] Khalid Bakhshaliyev, Muhammed Abdullah Canbaz, and Mehmet Hadi Gunes. Investigating characteristics of internet paths. *ACM Transactions on Modeling and Performance Evaluation of Computing Systems*, 4(3), August 2019.
- [26] Shkelzen Cakaj. The parameters comparison of the “starlink” leo satellites constellation for different orbital shells. *Frontiers in Communications and Networks*, 2, 2021.
- [27] Patricia Callejo, Marco Gramaglia, Rubén Cuevas, and Ángel Cuevas. A deep dive into the accuracy of ip geolocation databases and its impact on online advertising. *IEEE Transactions on Mobile Computing*, 22(8):4359–4373, 2023.
- [28] Massimo Candela, Enrico Gregori, Valerio Luconi, and Alessio Vecchio. Dissecting the speed-of-internet of middle east. In *Proceedings of IEEE INFOCOM 2019*, 2019.
- [29] Massimo Candela, Valerio Luconi, and Alessio Vecchio. Impact of the covid-19 pandemic on the internet latency: A large-scale study. *Computer Networks*, 182(9), December 2020.
- [30] Michael Chung. Diffusion Gaussian Kernel. *University of Wisconsin-Madison*, 2009.
- [31] Saahil Claypool, Jae Chung, and Mark Claypool. Comparison of tcp congestion control performance over a satellite network. In *Proceedings of International Conference on Passive and Active Network Measurement*, pages 499–512, 2021.
- [32] Lorenzo Corneo, Maximilian Eder, Nitinder Mohan, Aleksandr Zavodovski, Suzan Bayhan, Walter Wong, Per Gunningberg, Jussi Kangasharju, and Jörg Ott. Surrounded by the clouds: A comprehensive cloud reachability study. In *Proceedings of the Web Conference 2021*, 2021.
- [33] Jörg Deutschmann, Kai-Steffen Hielscher, and Reinhard German. Satellite Internet performance measurements. In *Proceedings of 2019 International Conference on Networked Systems (NetSys)*, 2019.
- [34] Sebastian Endres, Jorg Deutschmann, Kai-Steffen Hielscher, and Reinhard German. Performance of QUIC Implementations Over Geostationary Satellite Links. *arXiv preprint arXiv:2202.08228v2*, February 2022.
- [35] Nick Feamster and Jason Livingood. Measuring internet speed: Current challenges and future recommendations. *Commun. ACM*, 63(12):72–80, nov 2020.
- [36] Justus Fries, Trinh Viet Doan, Rohit Gupta, and Vaibhav Bajpai. An eight years perspective on the internet broadband infrastructure in the usa. In *Proceedings of IFIP Networking 2022*, 2022.
- [37] Lixin Gao. On Inferring Autonomous System Relationships in the Internet. *IEEE/ACM Transactions on Networking*, 9(6), 2001.
- [38] Google. Google developers - geolocation api. <https://developers.google.com/maps/documentation/geolocation/overview>, 2023.
- [39] Matthieu Gouel, Kevin Vermeulen, Olivier Fourmaux, Timur Friedman, and Robert Beverly. IP Geolocation Database Stability and Implications for Network Research. In *Network Traffic Measurement and Analysis Conference*, Online, United States, September 2021.
- [40] Jim Griner, John Border, Markku Kojo, Zach D. Shelby, and Gabriel Montenegro. Performance Enhancing Proxies Intended to Mitigate Link-Related Degradations. RFC 3135, June 2001.
- [41] Bingnan Hou, Changsheng Hou, Tongqing Zhou, Zhiping Cai, and Fang Liu. Detection and characterization of network anomalies in large-scale rtt time series. *IEEE Transactions on Network and Service Management*, 18(1), March 2021.
- [42] Hughes. Broadband Backup. Hughes website, 2023.
- [43] Geoff Huston. As names. <https://bgp.potaroo.net/cidr/autnums.html>, 2023.
- [44] ipinfo.io. ipinfo.io developers. <https://ipinfo.io/developers>, 2023.
- [45] Akshath Jain, Deepayan Patra, Peijing Xu, Justine Sherry, and Phillipa Gill. The ukrainian internet under attack: An ndt perspective. In *Proceedings of the 22nd ACM Internet Measurement Conference, IMC ’22*, page 166–178, New York, NY, USA, 2022. Association for Computing Machinery.
- [46] Ben Jones, Nick Feamster, Vern Paxson, Nicholas Weaver, and Mark Allman. Detecting dns root manipulation. In *Proceedings of the Passive and Active Measurement Conference 2016*, 2016.

- [47] jsDelivr. Open source cdn. how does it work? <https://www.jsdelivr.com/network/infographic>, 2023.
- [48] Mohamed M Kassem, Aravindh Raman, Diego Perino, and Nishanth Sastry. A browser-side view of starlink connectivity. In *Proceedings of the 22nd ACM Internet Measurement Conference*, pages 151–158, 2022.
- [49] Nicolas Kuhn, François Michel, Ludovic Thomas, Emmanuel Dubois, and Emmanuel Lochin. Quic: Opportunities and threats in satcom. In *2020 10th Advanced Satellite Multimedia Systems Conference and the 16th Signal Processing for Space Communications Workshop (ASMS/SPSC)*, pages 1–7, 2020.
- [50] L Lu. How to annotate AS relationships with geographic annotations? https://catalog.caida.org/recipe/how_to_annotate_as_relationships_with_geographic_annotations, 2016.
- [51] M- Lab. Privacy Policy. <https://www.measurementlab.net/privacy/>, 2023.
- [52] M-Lab. Annotation Service. <https://github.com/m-lab/annotation-service>, 2023.
- [53] Sami Ma, Yi Ching Chou, Haoyuan Zhao, Long Chen, Xiaoliang Ma, and Jiangchuan Liu. Network Characteristics of LEO Satellite Constellations: A Starlink-Based Measurement from End Users. *arXiv preprint arXiv:2212.13697*, December 2022.
- [54] Kyle MacMillan, Tarun Mangla, James Saxon, Nicole P. Marwell, and Nick Feamster. A comparative analysis of ookla speedtest and measurement labs network diagnostic test (ndt7). In *Abstract Proceedings of the 2023 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, SIGMETRICS ’23, page 41–42, New York, NY, USA, 2023. Association for Computing Machinery.
- [55] Gérard Maral, Michel Bousquet, and Zhili Sun. *Satellite communications systems: systems, techniques and technology*. John Wiley & Sons, 2020.
- [56] François Michel, Martino Trevisan, Danilo Giordano, and Olivier Bonaventure. A first look at Starlink performance. In *Proceedings of the 22nd ACM Internet Measurement Conference*, 2022.
- [57] Lai Yi Ohlsen and Stephen Soltesz. Most ndt clients migrated to ndt7. *MeasurementLab Blog*, June 2023.
- [58] Soohyun Park and Joongheon Kim. Trends in leo satellite handover algorithms. In *2021 Twelfth International Conference on Ubiquitous and Future Networks (ICUFN)*, pages 422–425, 2021.
- [59] Udit Paul, Jiamo Liu, Mengyang Gu, Arpit Gupta, and Elizabeth Belding. The importance of contextualization of crowdsourced active speed test measurements. In *Proceedings of the 22nd ACM Internet Measurement Conference*, IMC ’22, page 274–289, New York, NY, USA, 2022. Association for Computing Machinery.
- [60] PeeringDB. AS14593. <https://www.peeringdb.com/asn/14593>, 2023.
- [61] Daniel Perdices, Gianluca Perna, Martino Trevisan, Danilo Giordano, and Marco Mellia. When satellite is all you have: Watching the internet from 550 ms. In *Proceedings of the 22nd ACM Internet Measurement Conference*, pages 137–150, 2022.
- [62] Ingmar Poesse, Steve Uhlig, Mohamed Ali Kaafar, Benoit Donnet, and Bamba Gueye. Ip geolocation databases: Unreliable? *SIGCOMM Comput. Commun. Rev.*, 41(2):53–56, apr 2011.
- [63] Prolific. Prolific api for researchers (1.0). <https://docs.prolific.co/docs/api-docs/public/>, 2023.
- [64] Prolific. Prolific Researcher Help Center: Data Protection and Privacy. <https://researcher-help.prolific.co/hc/en-gb/articles/360009094594-Data-protection-and-privacy#heading-2>, 2023.
- [65] Prolific. Prolific’s participant pool: Its present and its future. <https://www.prolific.co/blog/prolifics-participant-pool-its-present-and-its-future>, 2023.
- [66] Will Reese. Nginx: The high-performance web server and reverse proxy. *Linux J.*, 2008(173), sep 2008.
- [67] RIPE NCC. RIPE Atlas legal information. <https://atlas.ripe.net/legal/>, 2023.

- [68] John P. Rula, Fabián E. Bustamante, and Moritz Steiner. Cell spotting: Studying the role of cellular networks in the internet. In *Proceedings of the 2017 Internet Measurement Conference, IMC '17*, page 191–204, New York, NY, USA, 2017. Association for Computing Machinery.
- [69] Luis A. Sanchez, Mark Allman, and Dr. Dan Glover. Enhancing TCP Over Satellite Channels using Standard Mechanisms. RFC 2488, January 1999.
- [70] SES. SES Exercises Call Option to Acquire 100% of O3b Networks. <https://www.ses.com/press-release/ses-exercises-call-option-acquire-100-o3b-networks>, 2016. Press Release.
- [71] Michael Sheetz. Investing in space: Is spacex’s starlink growing satellite internet market share, or taking it?, May 2023. <https://www.cnbc.com/2023/05/18/investing-in-space-where-spacexs-starlink-is-growing-market-share.html>.
- [72] Mahesh Sooriyabandara, Gorrry Fairhurst, Venkata Padmanabhan, and Hari Balakrishnan. TCP Performance Implications of Network Path Asymmetry. RFC 3449, December 2002.
- [73] Statcounter. Desktop browser market share worldwide. <https://gs.statcounter.com/browser-market-share/desktop/worldwide>, 2023.
- [74] Srikanth Sundaresan, Xiaohong Deng, Yun Feng, Danny Lee, and Amogh Dhamdhere. Challenges in inferring internet congestion using throughput measurements. In *Proceedings of the 2017 Internet Measurement Conference, IMC '17*, page 43–56, New York, NY, USA, 2017. Association for Computing Machinery.
- [75] Matteo Varvello, Hyunseok Chang, and Yasir Zaki. Performance characterization of videoconferencing in the wild. In *Proceedings of the 22nd ACM Internet Measurement Conference*, pages 261–273, 2022.
- [76] Viasat. Backup Internet. Viasat website, 2023.
- [77] ViaSat. Viasat Secure Network Accelerator. <https://eguide.field.viasat.com/wp-content/uploads/SNA-User-Guide.pdf>, 2023.
- [78] Wondernetwork. Global Ping Statistics → Manila and Tokyo. <https://wondernetwork.com/pings/Manila/Tokyo>, 2023.
- [79] Maya Ziv, Liz Izhikevich, Kimberly Ruth, Katherine Izhikevich, and Zakir Durumeric. ASdb: A system for classifying owners of autonomous systems. In *Proceedings of the 21st ACM Internet Measurement Conference*, 2021.
- [80] Zvleo. Market-leading url database and web content categorization services. <https://zvleo.com/zvelodb-url-database/>, 2023.

Appendix

A Ethics

Our work uses the performance traces from *(i)* two widely-used speed test measurement suites and testbeds (operated by M-Lab and RIPE) and *(ii)* human participants through Prolific. The data from former sources are anonymized [51, 67] and are publicly available (through Google BigQuery and APIs). In fact, M-Lab recommends to use the data for research by waiving the copyright and related rights to the data. In case of latter, Prolific anonymises all the participant information [64]. There are no methods available to decode any Personal Identifying Information from Prolific. We also contacted the Institutional Review Board (IRB) office at our University and they deemed that this was not human subjects research and as the tool is mainly automated to run and collect results.

B Exhaustive List of SNOs

Table 3 shows a curated list of 67 ASNs belonging to 41 SNOs we have identified using the methodology we developed in this paper, specifically “ASN-to-SNO Mapping” described in Section 3.2.

C BGP Peering Analysis of SNOs

Our BGP-peering-based PoP characterization described in Section 4 allows us to examine and compare different SNOs’ ground infrastructures as well as their historical

SNO	ASN	SNO	ASN	SNO	ASN
arqiva	15641	isotropic	36426	sound & cellular	63215
avanti	39356	kacific	135409	speedcast	38456
awv	46869	kvh	25687	ssi	22684
colinanet	262168	lepton (kymeta)	20304	starlink	14593
comsat	36614	linkexpress	20660		27277
comsat (png)	136940	marlink	5377		
comtech	394318		44933	telalaska	10538
elara	262927		55784	telesat	19036
eutelsat	204276		8841	televera	265515
	34444		210314	thaicom	63951
	15829		8264	ultisat	393439
globalsat	28503		37101	viasat	13955
gravity	131202		393938		25222
hellas-sat	41697	maxar	203101		46536
hughes	28613	navarino	133933		18570
	1358	netsat	1821		16491
	63062	network innovations	395786		40306
	12440	nomad global	60725		7155
	44795	o3b	800		40310
	6621	oneweb	64294		40311
intelsat	26243	panasonic	201554		23354
	46982	ses	12684		31515
io	17411			worldlink	11902

Table 3: Curated list of SNO-ASN mapping captured from ASdb and manually visiting each operator website.

evolution over time. Figure 12 shows BGP peering visualization of the SNOs evaluated in Figure 4c, as captured by BGP route-views collected in January, 2023. Figure 13 compares representative growth patterns of several SNOs. For this analysis, we use BGP route-views snapshots captured from three separate periods with one year interval (from 2021/1 to 2023/1).

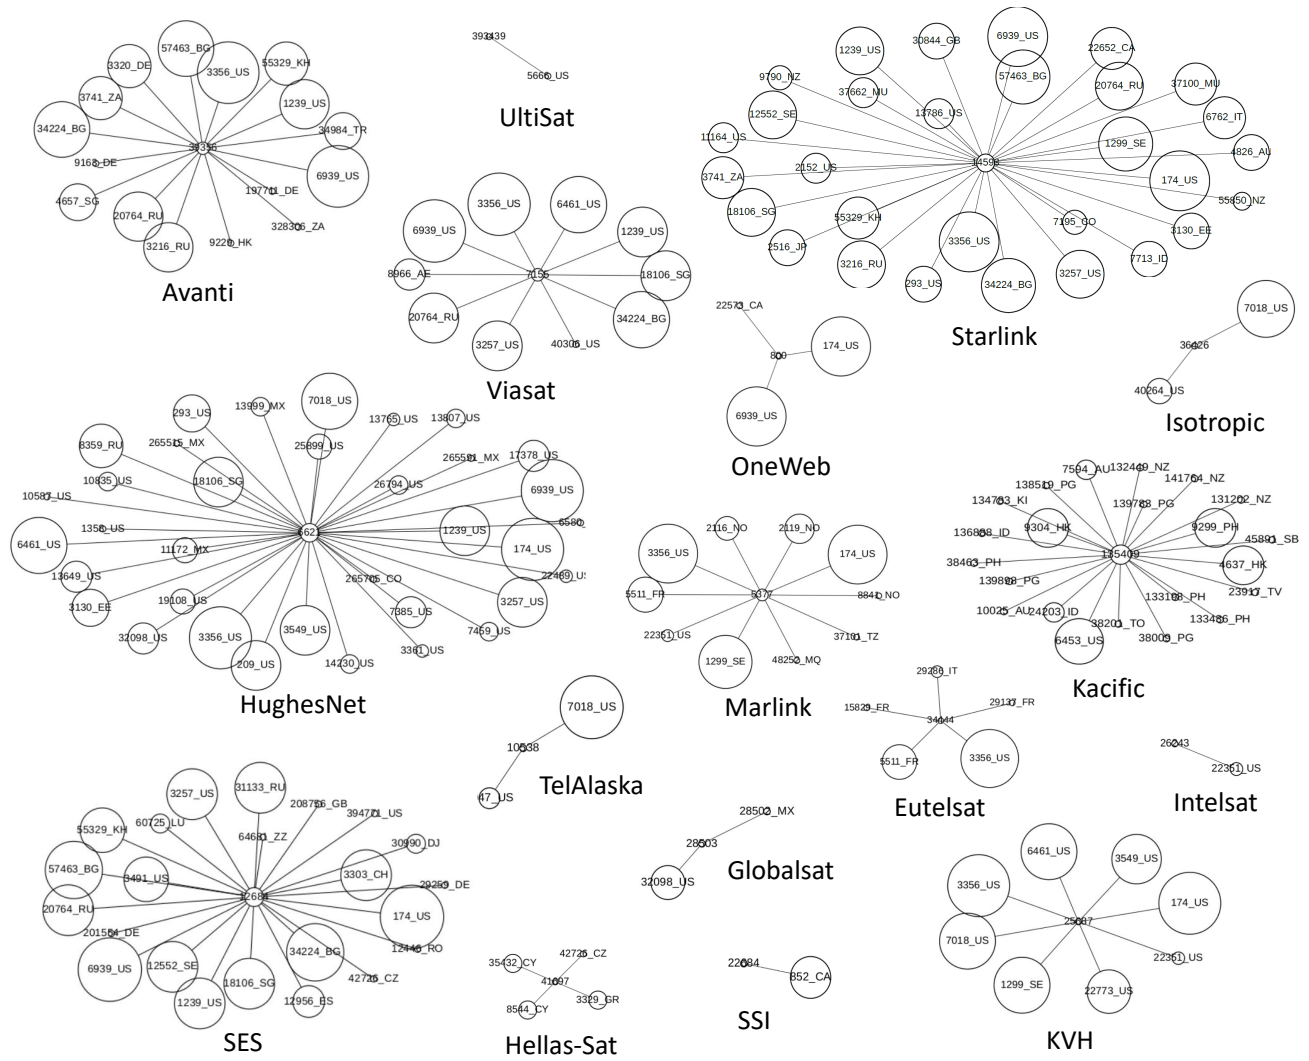


Figure 12: BGP peering visualization of SNOs (based on BGP route-views collected on 2023/1/1).

D Prolific Census Results

Figure 14 visualizes feedback (from “very poor” to “very good”) from 56 Prolific testers who are also subscribers of HughesNet, Starlink, and Viasat.

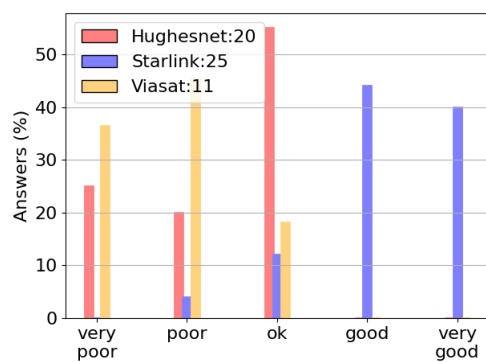


Figure 14: Feedback from 56 Prolific testers on the quality of their SNO. Scores from 1 (very poor) to 5 (very good). Data collected in March 2023.