Understanding DNS Dynamics over the Starlink Network

Robert Richter
Hasso Plattner Institute, *University of Potsdam*
robert.richter@hpi.de

Vaibhav Bajpai
Hasso Plattner Institute, *University of Potsdam*
vaibhav.bajpai@hpi.de

***Abstract*—Low-Earth Orbit (LEO) satellite networks have emerged as a viable Internet connectivity option, yet their Domain Name System (DNS) behavior remains largely unexplored. This paper presents the first dedicated measurement study of DNS performance in the Starlink network. Using RIPE Atlas infrastructure, we analyze DNS resolution times across various countries over a 13-month observation period from January 2025 to February 2026. We investigate three research questions: which DNS resolvers Starlink users commonly employ, how different resolvers compare in performance, and what influences DNS latencies. Our measurements reveal that DNS latency in Starlink is predominantly determined by the satellite hop rather than terrestrial routing, with the CGNAT hop (`100.64.0.1`) accounting for approximately 69 % of total latency in traceroute measurements. Reducing this latency would only work by locating the resolver closer to the Starlink user. We also examine Starlink's subscription-based DNS handling and find that unverified users are assigned a restricted resolver (`34.145.127.1`) that resolves only specific domains. Additionally, we identify SpaceX-operated IPv6 DNS resolvers that appear co-located with PoPs, potentially offering performance benefits through geographic proximity. Our findings provide insights into the design of the DNS ecosystem for LEO satellite networks.**

***Index Terms*—LEO Satellite Networks, Starlink, DNS, Resolvers, Performance, Internet Measurements**

I. Introduction

LEO satellite networks have transformed global Internet connectivity. Unlike traditional GEO satellites orbiting at approximately 35 786 km, LEO constellations operate at altitudes between 500–1200 km, reducing communication latency and enabling interactive applications previously infeasible over satellite links. Starlink, operated by SpaceX, has emerged as a well-established provider with over 10 000 active satellites. The rapid deployment of LEO constellations has extended broadband access to remote regions, maritime vessels, and aircraft, serving millions of users who previously lacked reliable connectivity options.

The maturation of LEO satellite networks makes comprehensive performance analysis both timely and essential. Starlink has transitioned from an experimental service to critical infrastructure for many users, yet our understanding of its network behavior remains incomplete. While recent studies have examined general latency characteristics and content delivery performance, the DNS has received little attention in the LEO context. Given the Internet's reliance on DNS resolution, understanding its performance characteristics and behavior is crucial for both end users and network operators.

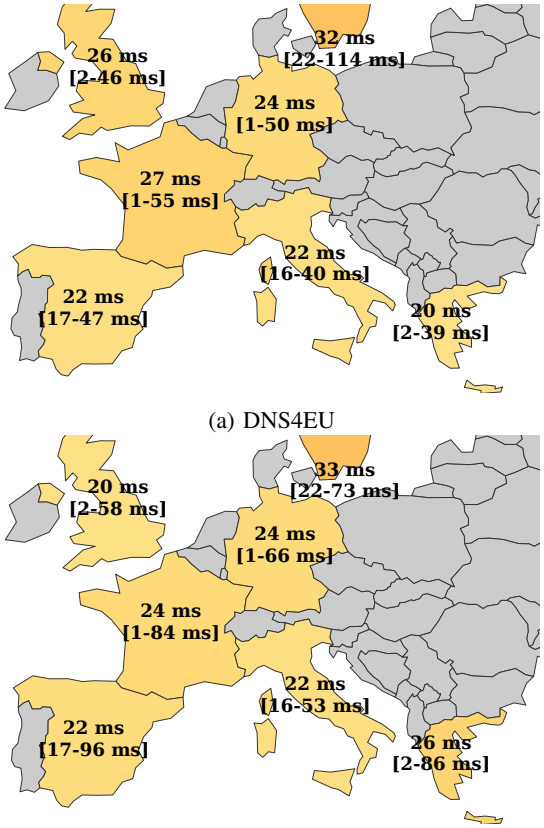

(a) DNS4EU

(b) Quad9

Fig. 1: Starlink latencies of Quad9 and DNS4EU in various European countries. The top number is the median latency, while the numbers below show the 5th to 95th percentile of latencies using data from local measurements.

DNS performance in LEO networks presents unique challenges that distinguish it from terrestrial environments. In conventional networks, geographic proximity between clients and resolvers typically correlates with low latency, enabling straightforward optimization through anycast deployment [1] or CDN networks. LEO satellite networks disrupt this assumption entirely. Packets from a user terminal must traverse the satellite constellation to reach a ground station, then route through a PoP before reaching a receiver in the terrestrial In-

ISBN 978-3-903176-82-9

ternet. This architecture introduces variable latency (as shown in Figure 1) that depends, among others, on satellite positions, inter-satellite links, and ground station selection. These factors complicate traditional DNS optimization strategies and raise questions about resolver placement, caching effectiveness, and protocol selection. Therefore, we outline the following RQs:

RQ 1 How does Starlink set default resolvers?
RQ 2 What are major influences on DNS latency?
RQ 3 How do different resolvers perform for Starlink users? Is there a recommended resolver?
RQ 4 Can we observe differences in performance when varying the DNS protocol and query types?

To address the four RQs, we leverage the RIPE Atlas built-in measurement infrastructure to collect DNS resolution data from Starlink-connected probes over a 1-year period across various countries. Our findings are:

1. **Impact of DNS Resolvers** (Section V): We investigated the performance of various DNS resolvers, including SpaceX-operated resolvers, public DNS services (Quad9, DNS4EU, Google, Cloudflare), and local resolvers, identifying SpaceX-operated resolvers as the fastest due to their proximity to PoPs.

2. **Impact of Caching and Query Types** (Section VII): We analyzed the effects of caching and query types (A vs. AAAA) on DNS resolution times, finding that caching significantly reduces latency, while query types have a negligible impact.

3. **Impact of CGNAT and Satellite Hops** (Section VI): DNS latency in Starlink networks is predominantly determined by the satellite hop, with the CGNAT hop accounting for approximately 69 % of total latency.

4. **Impact of Encrypted DNS** (Section VIII): We compared the performance of traditional DNS (Do53 [2], DoTCP [3]) to encrypted DNS protocols (DoT3 [4], DoH2 [5]) and found that encrypted DNS protocols (DoE) incur significant latency overhead due to additional round-trip and encryption-related delays.

We also discovered that unverified Starlink users are assigned a restricted resolver (`34.145.127.1`) that resolves only specific domains (*e. g.*, *starlink.com* and *spacex.com*), enabling limited connectivity for subscription activation.

Further, we define custom RIPE Atlas measurements to obtain additional data on specific DNS resolvers. The measurements are enhanced with custom measurements from our self-operated measurement device. The measurements are shown in Figure 3. Details on the methodology are outlined in Section IV. In the following, we address the first RQ in Section II and Section V-A by elaborating on the existing documentation and further filling information gaps. The second RQ is examined in Section VI using built-in measurements from RIPE Atlas. As those measurements only consider root servers, in Section V-B we conduct measurements on specific DNS resolvers to gain insights into RQ 3. Last, in Section VII, we use our custom measurements to answer RQ 4.

This work contributes to an understanding of the Starlink DNS infrastructure (Sections V-A and VI). It visits various common DNS resolvers and their latencies (Section V). We identify important aspects of LEO satellite networking that must be considered for future infrastructure and research efforts.

All of our code and datasets are published via GitHub[1].

II. Background

Over the last few years, LEO satellite networking has become prominent. Various SNOs deployed their NSCs, with Starlink having the most satellites by far (see Figure 3). Other noteworthy NSCs are Amazon Leo (formerly Project Kuiper), which recently deployed a larger number of satellites, and OneWeb, the only European service with a notable number of satellites.

Networking in LEO comes with specific challenges. Sending packets passes through various intermediate nodes, such as a user antenna, various satellites, a GS, and a PoP. That leads to problems when using traditional networking technologies. For instance, the argument that geographic proximity usually implies low latency does not work (*e. g.*, as shown in studies about CDNs [6], [7]) and packet loss is high [8], [9].

The DNS ecosystem in LEO NSCs has seen only little attention to date. Additionally, for Starlink, little is known about the system's default resolver choice. The company has documented that a user without a verified Starlink subscription defaults to the resolvers `34.145.127.1` [10]. After a successful verification, the resolver is updated to a SpaceX-operated DNS recursive resolver. We found that those resolvers were not publicly available. Users can further choose to change the default resolver to a public one (*e. g.*, `1.1.1.1` or `8.8.8.8`). While we cannot fully verify the process, we believe the process of default Starlink DNS resolver choice is close to the one illustrated in Figure 2.

In theory, using a SpaceX-operated resolver should provide a practical advantage if the resolver is colocated with the Starlink PoP. Geographical proximity between PoP and ISP resolver may reduce terrestrial latency overhead to less than a millisecond as the resolver is on the data path and does not introduce additional intermediate hops. The PoP per user is rarely changed [11], so the resolver can even be predetermined. Other properties (*e. g.*, DDR [12] or ECH [13] support) of the resolver itself remain future work and will not be part of this study.

III. Related Work

In the previous years, a number of performance measurements have been conducted [14], [9], [15], [11], [8], [16], [17]. Research found that Starlink is capable of providing low-latency Internet connectivity as low as 20–30 ms [11]. However, this comes with high packet loss [8] and high geographic variation [9]. Geographic proximity does not correlate with low latency anymore, which becomes apparent by observing CDNs [6], [17], [18], [19]. Bose et al. find *"fundamental architectural misalignments between satellite-based Internet access and terrestrial content delivery infrastructure"* [19]. Public DNS services have properties similar to those of CDNs

[1] See https://github.com/diic-starlink/Understanding-DNS-Dynamics-over-the-Starlink-Network

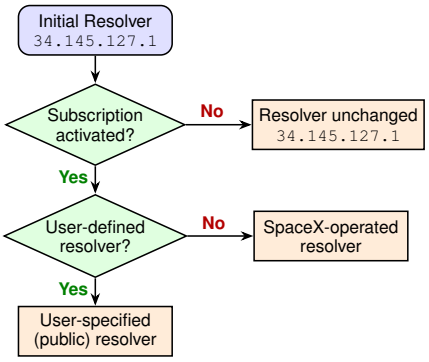

Fig. 2: Assumed DNS resolver setting for Starlink users

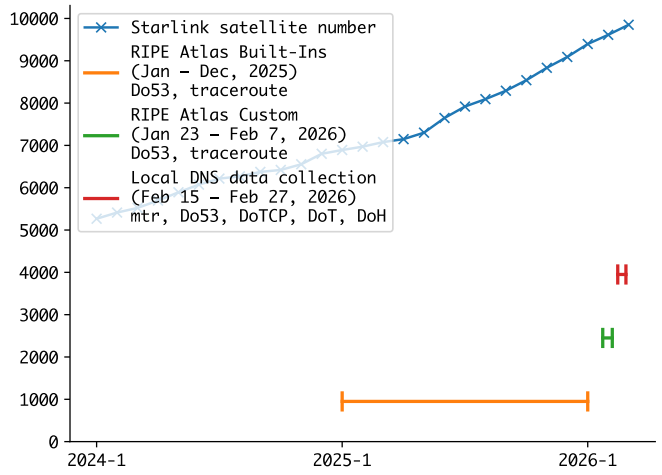

Fig. 3: Number of satellites in the Starlink constellation and the timespans of our measurements

(*e. g.*, global anycast infrastructure). Therefore, studying DNS is important. A part of the CDN studies also examines DNS performance for Google, Cloudflare, and Quad9 resolvers. They find that the user terminal to PoP latency accounts for approximately 75.82 % of the latency (averaged among all studied vantage points). Further, they find DNS latencies for common public resolvers to be around 20–40 ms, which aligns with our findings for Do53 (see Section IV-C). DoE, which is not yet studied, may show very different results because of protocol deficiencies in LEO satellite networks [17], [16].

IV. Methodology

This section describes our measurement methodology for analyzing DNS characteristics in Starlink networks (*i. e.*, AS14593). We employ RIPE Atlas built-in measurements, custom active measurements for specific DNS resolvers, and verification using a Starlink-connected measurement device on the university's campus. The conducted measurements are listed in Figure 3.

All measurement artifacts and code will be publicly released after the review process.

A. Global LEO DNS Measurements using RIPE Atlas

We leverage RIPE Atlas to obtain geographically distributed DNS measurement data from Starlink-connected probes. Our approach combines two complementary data sources. RIPE Atlas probes continuously perform built-in DNS and traceroute measurements to root DNS servers as part of the platform's standard operation. Those measurements have IDs 50XX and 100XX, depending on the targeted root server. The DNS queries use Do53. We collect these measurements from all probes that have ever delivered built-in measurement results during our observation period, regardless of their current connection status, as long as they gave measurement results in the timeframe. This inclusive approach maximizes our dataset coverage and enables analysis of probes with intermittent connectivity.

For targeted resolver performance analysis, we also conduct custom RIPE Atlas measurements from all active Starlink-connected probes. Each measurement queries the domain `<anonymized>` using DNS record type A over IPv4. We target the following DNS resolvers:

- **Quad9** (`9.9.9.9`): A widely-deployed public recursive resolver operated by the Quad9 Foundation, providing a baseline for anycast DNS performance.
- **DNS4EU Unfiltered** (`86.54.11.100`): A public Europe-centric DNS resolver initiative by the European Commission. Offers a couple of resolvers for different goals (*e. g.*, child protection or ad blocking).
- **Local Resolver**: The default resolver configured on each probe, capturing the characteristics default DNS infrastructure or user-configured alternatives.

The resolvers `1.1.1.1` (Google's DNS) and `8.8.8.8` (Cloudflare's DNS) were omitted during our custom measurements as RIPE Atlas only allows 25 concurrent measurements to the same target [20]. We were not able to schedule a measurement without violating this rule. However, some users specify public resolvers (*e. g.*, Google or Cloudflare) as their local resolvers, allowing us to analyze data for those resolvers even if we were unable to explicitly start measurements with them.

B. Local DNS Measurements

At the university's campus, we run local measurements to complement the RIPE Atlas measurements. We use a Starlink-connected VM that runs DNS measurements against Google DNS, Cloudflare DNS, Quad9, and DNS4EU. We use the `kdig` tool to measure DNS performance. Specifically, we vary (1) query type (A and AAAA), (2) protocol[2] (Do53, DoTCP, DoT3, and DoH2[3]), and (3) cache state (cached vs. uncached). For cached queries, we query `google.com`. For uncached queries, we generate a unique new DNS entry before each query. The generated domain has the format `{UUIDv4}.{SLD}`. The corresponding RR is either A or

[2]Please note that further encrypted DNS protocols exist (*e. g.*, DoQ or DoH3), but were not studied due to missing resolver or `kdig` support.

[3]`kdig` uses HTTP/2, as it uses `libnghttp2`.

AAAA and points to localhost (*i. e.*, `127.0.0.1` or `::1`). After the DNS query, the entry is deleted again.

The experiment yields a three-tuple of measurement inputs: (*RR*, *protocol*, *resolver*). Therefore, each measurement collects 16 data points. The experiment is repeated every 30 min. Assuming no measurement fails, we gather 768 data points per day. The DNS measurement run is followed by an `mtr` measurement to the respective resolver.

C. Querying the `34.145.127.1` Resolver

We wanted to determine which addresses the SpaceX-operated resolver `34.145.127.1` returns meaningful responses for. Meaningful means a response that can actually be used for further communication (*e. g.*, an IPv4 address for A queries). For most addresses, the resolver returns an empty response without an error, but with a warning that the resolver does not support recursion. We assume that the resolver maintains a predefined list of domains it resolves, whereas all other domains do not. To determine a list of SLDs, we query the resolver for common DNS RRs for each entry. The entries originate from the Tranco 100k List (ID 2N6G9; January 2026). Querying happens in the following order: `A`, `AAAA`, `TXT`, and `CNAME`. If an entry returns a valid entry, no further RRs of the address are queried. For ethical crawling, we set a timeout of 200 ms between each query of a RR and 1 s between each domain address.

V. Impact of DNS Resolvers

First, we aim to understand the DNS ecosystem in Starlink. Repeating, Starlink likely uses the default resolver choice based on the decision graph illustrated in Figure 2. Therefore, we study the ecosystem without an active subscription (Section V-A) and following DNS performance for different resolvers (Section V-B).

A. DNS Resolver Choice without a Starlink Subscription

Most users likely already have a working subscription when they install their Starlink antenna, so the behavior without an active subscription will not apply. However, connectivity remains active even when a subscription is not present. The assumed purpose is to enable users to acquire a Starlink subscription when they already have a Starlink antenna and no other means of Internet connectivity. When we first set up a Starlink Mini antenna, we received an hour of free service even though a subscription wasn't enabled (as noted in the Starlink mobile app). We were able to regularly access the Internet during this time. After the hour passed, we were only able to access `starlink.com` and `spacex.com`.

As documented [10], the DNS resolver `34.145.127.1` is set for users who have not been verified by the Starlink network (*i. e.*, have no active subscription). A lookup using IPinfo shows that the DNS resolver `34.145.127.1` is hosted within Google's network in the US. When testing the resolver with different domains, it becomes apparent that it does not respond to all domains, including major ones like `google.com` or `cloudflare.com`. We followed up to determine which domains resolve on that resolver. We tried resolving the Tranco Top 100k SLD Domains [21]. For a description of the exact querying strategy, see Section IV-C. Out of all 100k domains, only `starlink.com` and `spacex.com` resolved to an IP address. Also, neither website attempted to load 3rd-party assets that do not require any additional resolvable domains. A minor `mtr` measurement from the university campus via Starlink revealed an average ping RTT of ≈ 161 ms. While this is a high latency, it is also not surprising, as the resolver is not intended for active low-latency connectivity.

In summary, Starlink's subscription centers on setting the DNS resolver to one that resolves only for specific domains, enabling users to re-enter a subscription without using the full Internet. Manually setting a different DNS resolver does not yield further results. Likely, Starlink uses label switching to prevent packets from being forwarded when a specific IP address is detected.

After subscribing to Starlink, connectivity is fully available. According to Figure 2, the DNS resolver is set to a SpaceX-operated resolver if the user does not specify other preferences. More specifically, the user either chooses a public DNS service or will default to a SpaceX-operated resolver. During our RIPE Atlas measurements (outlined in Section IV-A), we encountered a couple of IPv6-capable SpaceX-operated DNS resolvers. We were not able to connect to those addresses. It is very likely that those resolvers are co-located with PoPs to reduce the length of the data path, while being inaccessible by non-Starlink users to preserve resources. Because the data path is significantly shorter, we expect SpaceX-operated services to be faster than public DNS resolvers. Still, resolving resources, caching, and recursive resolution may yield differences, but we do not see signs that public services are advantageous.

In conclusion, comparing DNS latencies across resolvers is largely about measuring the PoP to resolver latency, as packets follow the same data path before reaching the resolver, regardless of the resolver's properties. However, if the user antenna to PoP latency dominates the RTT, the choice of resolver may play a negligible role (as long as they provide the same features).

Users may decide to specify a custom DNS service using the Starlink mobile app. Common options are Cloudflare DNS, Google DNS, and Quad9. In January 2023, DNS4EU also started offering public services. Using RIPE Atlas and the university's Starlink connection, we measure the performance of those four public DNS services.

In Section V-B, we look at the results of measurements taken on RIPE Atlas. Those measurements target Quad9, DNS4EU, and the probe's locally configured resolver. We also compare the results with our own measurements. Following Section VII, we examine in detail the comparison of individual protocols, caching behavior, and IP version (*i. e.*, A vs AAAA queries).

B. Studying the Resolvers from Global Vantage Points

We use RIPE Atlas to run measurements using resolvers from DNS4EU, Quad9, and the local resolver. The advantage

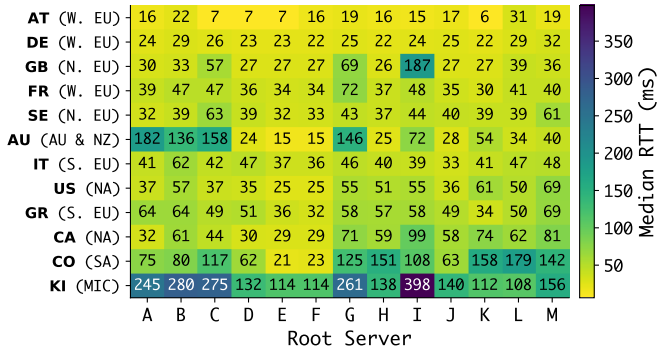

(a) January 2025

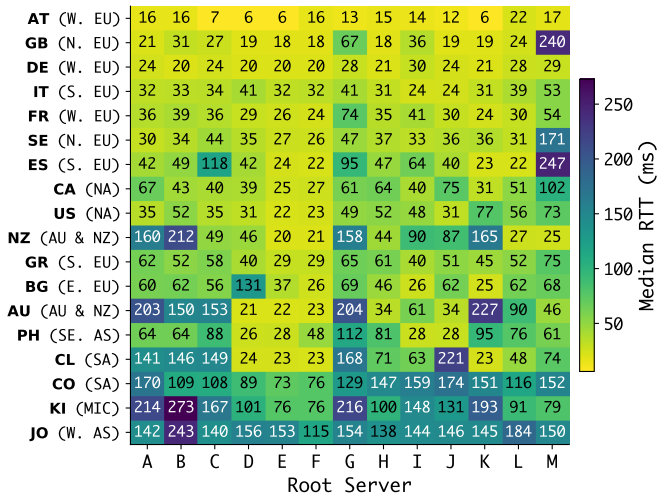

(b) December 2025

Fig. 4: Median DNS latencies to root servers in the first and last month of 2025 using RIPE Atlas built-in DNS measurements. Latencies are sorted from top to bottom.

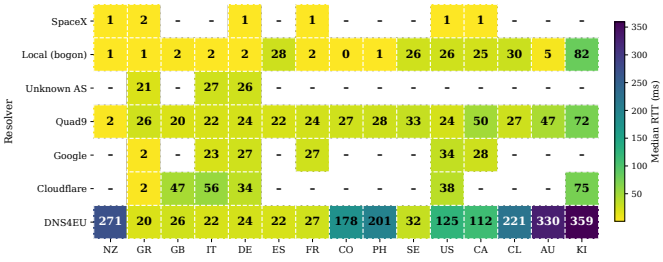

Fig. 5: Median DNS latency by resolver and country using data from RIPE Atlas custom measurements.

over the measurements in Section VI is that we catch the performance of recursive resolution. However, we cannot observe caching effects, as measurements in RIPE Atlas cannot point to `google.com` (as explained in Section IV-A). Other domains do not offer sufficient certainty that they will be cached. Further, RIPE Atlas only offers limited DNS protocol possibilities. Therefore, we analyze Do53 in this section and elaborate on protocols in Section VII. Figure 5 illustrates the

TABLE I: CGNAT hop latency portion of total latency using data from traceroute measurements in custom RIPE Atlas experiments.

Country	Ratio
Kiribati	76.75 %
France	73.86 %
Canada	79.94 %
Austria	111.99 %
Philippines	13.05 %
United States	71.47 %
Germany	102.02 %

observed DNS latencies, including Quad9, DNS4EU, Cloudflare, Google, and SpaceX resolvers. Most data is available for Quad9 and DNS4EU, as these resolvers were explicitly defined in the measurements. Other resolvers are the result of local resolver configurations on RIPE Atlas probes. The observed local resolver address was mapped to an organization using IPinfo, yielding data for Cloudflare, Google, and SpaceX resolvers. Some addresses were also bogon addresses that we were not able to map to a specific resolver. Further, we observed non-bogon addresses, but IPinfo could not find the corresponding AS (categorized as *Unknown AS*).

Comparing Quad9 and DNS4EU, Quad9 has significantly lower latency globally. However, focusing on European countries, DNS4EU shows lower latency and less variation. That aligns with DNS4EU's aim of serving European countries, whereas Quad9 does not focus on a specific geographic region. Google offers similar latency compared to DNS4EU. More specifically, we observe that SpaceX resolvers have the lowest latencies. As expected, this may be due to a shorter data path resulting from assumed colocation with PoPs. We can conclude that SpaceX-operated resolvers exhibit lower latency than public DNS services, but data on caching effects and protocol support is missing. It is not clear whether SpaceX-operated resolvers support common DoE protocols and whether they support supplemental features such as DDR. Those resolvers are not publicly accessible, which requires a study to implement a detection method conducted from existing Starlink terminals.

VI. Impact of CGNAT and Satellite Hops

We used built-in RIPE Atlas measurements to determine baseline DNS latencies (*i. e.*, measurement IDs 100XX). The measurements use probes that are or were connected to RIPE Atlas. Additionally, we only considered countries with at least two probes to ensure data confidence. It becomes apparent that there is a major difference between countries, as also claimed by related work [8], [9], [11]. Latency differences in DNS measurements are typically attributed to caching behavior [22] (and the corresponding TTL), but also other infrastructure-dependent factors (*e. g.*, distance and congestion). We observe similar behavior, but the results vary based on the target server. Figure 4 shows our results on DNS latency to all root servers across various countries in the first and last month of 2025. Most DNS queries resolve in less than 50 ms. We

also observe a right-skewed distribution (*i. e.*, many countries have low median DNS latencies, while some show clearly higher ones). Additionally, some countries exhibit significantly higher variation than others (*e. g.*, Kiribati and Jordan). In general, European countries show the lowest latencies. Either the European Starlink infrastructure is well expanded, or the placement of the root servers is better. However, we can derive that Starlink clearly focuses on specific regions (if intended or not). Further, we observe that some root servers exhibit better results than others. For instance, *F* (366 sites [23] worldwide) achieves less than 30 ms in most countries. Other root servers, such as *M* (29 sites worldwide), struggle to maintain latencies below 50 ms.

Looking at the time series, we see more stable latencies in January 2025 than in December 2025. We observe this for all countries where data was available in both timeframes (*e. g.*, see Colombia). It seems highly correlated with the deployment sites of root servers. Root servers with many sites are likely to be close to a PoP, and latency will be low. Another explanation could be congestion in the Starlink network, which could lead to variance, but we have no reliable data to support that.

Furthermore, we study the causes of major latency variation to identify its sources. Generally, we assume latency variation is due to inefficient PoP identification per probe (*i. e.*, a suboptimal PoP is chosen for data transmission) or to distant root servers. The latter is unlikely, as the chosen root servers are highly available worldwide (see root server map [23]), while suboptimal PoP choice has been observed in earlier work [7], [6]. Please note that it is very unlikely that DNS caching is a source of latency variation as the measurements run against root servers. To find the source of variation in Starlink's DNS latency, we analyzed RIPE Atlas built-in traceroute measurements (*i. e.*, measurements with ID 50xx). We looked at the traceroute results for individual probes. We observe that in most traceroutes, IPv4 `100.64.0.1` appears and causes a major latency spike. Starlink documents [10] this address as the private IP address (from prefix `100.64.0.0/10`) assigned to users for CGNAT. This hop induces the highest latency increase in most of our traceroute measurements. Our traceroute latencies to excerpts of countries are shown in Table I. They compare the hop's latency to `100.64.0.0/10` with that of the total latency. As traceroute is not a precise latency measurement, we also observe percentages above 100 %. Among the studied countries, the ratio amounted to approximately 68.83 %. Please note that Table I does not show all studied countries. That implies that most of the variation in Starlink DNS measurements stems from the hops toward CGNAT. To further confirm this observation, we visualized the hop to `100.64.0.1` in Figure 7 and compared it against the total latency. We find that in Germany, the "CGNAT-hop" accounts for nearly all of the latency, whereas in the Philippines, Jordan, and Colombia, the later hops also significantly contribute to the total latency. We assume that the observed latency is not caused by CGNAT, but by the satellite constellation. Packets pass first through the satellite constellation, and when they reach a PoP, CGNAT is applied. However, our traceroute

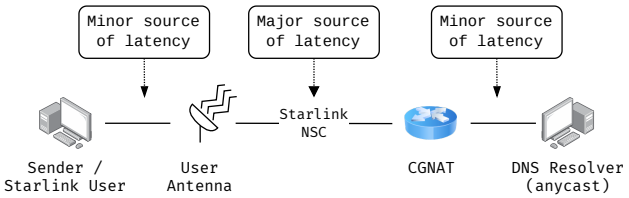

Fig. 6: Illustration of the DNS latency influences as observed by traceroute measurements.

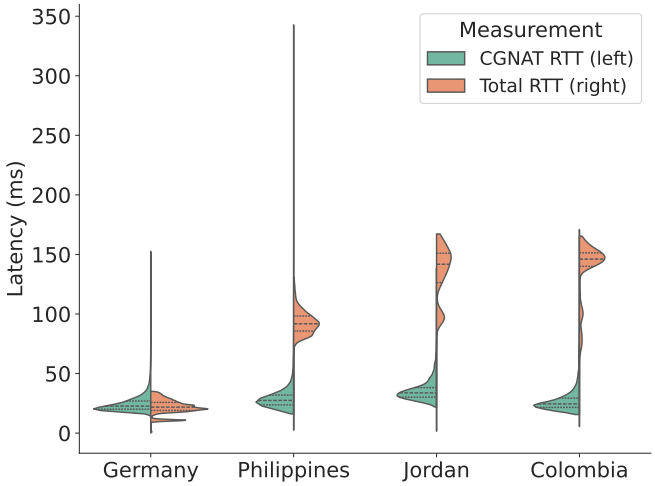

Fig. 7: Comparison of total traceroute latency and latency to hop `100.64.0.1` from RIPE Atlas built-in traceroute measurements in various countries. Capped at the 90th percentile.

measurements do not capture the satellite constellation, which operates below the IP layer. Therefore, we assume that the larger latency hop is due to routing through multiple satellites and other unknown intermediate nodes.

In summary, reducing DNS latency is not attributable to the resolver infrastructure, provided a global anycast infrastructure is in place and latency is generally low. For the most part, it depends on Starlink's infrastructure. Studying DNS latencies should focus on antenna to PoP latency, when measurements show public resolvers have similar low latency.

Takeaway 1. Latency to root servers is largely determined by the antenna to PoP latency ($\approx 68.83\,\%$ of total latency with differences by country). We can conclude that the user antenna to PoP DNS latency will dominate the total latency. Figure 6 illustrates this finding. Server-side improvement are unlikely to improve performance.

VII. Impact of Caching and Query Types

We consider multiple dimensions of DNS performance in our local measurements. Broad median latencies are illustrated in Figure 8. We study performance characteristics related to caching, protocol choice, and RRs. We hypothesize that,

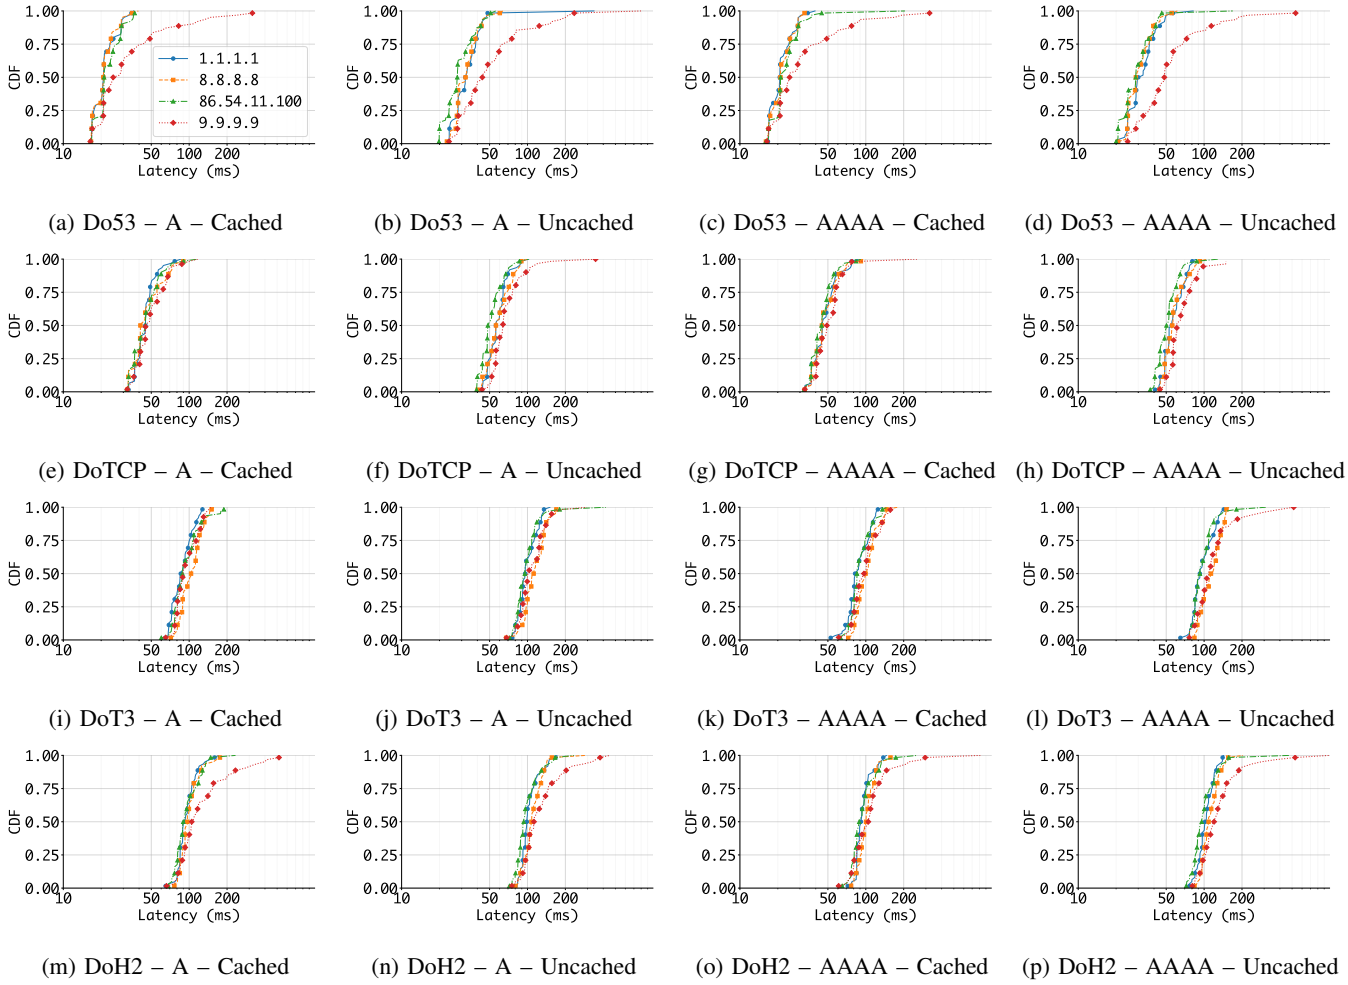

(a) Do53 – A – Cached (b) Do53 – A – Uncached (c) Do53 – AAAA – Cached (d) Do53 – AAAA – Uncached

(e) DoTCP – A – Cached (f) DoTCP – A – Uncached (g) DoTCP – AAAA – Cached (h) DoTCP – AAAA – Uncached

(i) DoT3 – A – Cached (j) DoT3 – A – Uncached (k) DoT3 – AAAA – Cached (l) DoT3 – AAAA – Uncached

(m) DoH2 – A – Cached (n) DoH2 – A – Uncached (o) DoH2 – AAAA – Cached (p) DoH2 – AAAA – Uncached

Fig. 8: CDFs of DNS resolution times using various protocols, resource records, and caching behavior.

in the context of Starlink, caching does not play a larger role than in traditional terrestrial networks. Furthermore, we hypothesize that querying different RRs does not lead to systematic differences. Protocol-wise, we may see substantial differences because multiple round trips may be involved (*i. e.*, all protocols except Do53). This is a key factor because, as shown in Section VI, traversal through the satellite constellation dominates latency, especially for geographically close receivers.

First, we study the influence of caching across the four resolvers. Since Do53 is the most established DNS protocol [24], [25], [26], we begin by analyzing caching effects for Do53. The results are shown in Figures 8a and 8b and Figures 8c and 8d. As expected, cached domains resolve faster than uncached ones. The minimum latency for cached domains is $16\,\mathrm{ms}$, while the minimum latency for an uncached domain is $19.8\,\mathrm{ms}$. Similarly, the median cached latency is $\approx 21\,\mathrm{ms}$, while for uncached domains it is $\approx 30\,\mathrm{ms}$. Interestingly, Quad9 shows a very heavy-tailed distribution, unlike the other resolvers. Cloudflare, Google, and DNS4EU show very similar latencies. Overall, these caching effects are consistent with related work in terrestrial networks [22]. Next, we take a closer look at query types. Our results are illustrated in Figures 8a and 8c for the cached domain and in Figures 8b and 8d for uncached domains. Neither for A nor AAAA records do we observe substantial differences. Both RRs resolve with similar latency. Quad9 is the only resolver to show a heavy-tailed distribution, and it does so consistently across all four combinations.

Takeaway 2. We conclude that A and AAAA queries exhibit similar latency across the studied resolvers. Furthermore, caching has a significant effect on latency, but not to a different extent than in the terrestrial Internet.

VIII. Impact of DNS over Encryption

Next, we study the effect of different DNS protocols on latency. In general, a user may use unencrypted DNS (Do53), or encrypted DNS (*i. e.*, DoH2 and DoT3). Encrypted DNS is expected to have higher latency compared to Do53, as already observed in Figure 8. In Starlink, multiple RTTs may cause significant problems due to unpredictable latency

overhead. Additionally, high packet loss [8], [11], [27] further complicates the communication. As we found before, A and AAAA queries do not show different behavior. Therefore, for simplicity, we focus on A queries. First, we study the effects of querying a cached domain. We observe that Do53 has the lowest latencies among all measured dimensions. As illustrated in Figure 8, cached DoTCP is approximately two times the latency of Do53, which aligns with the extra round trip introduced by TCP (*i. e.*, round trip from TCP three-way handshake). Encrypted DNS is approximately four to five times slower compared to Do53, with no significant difference in DoH2 and DoT3. Both DoH2 and DoT3 require three round-trips when TLS 1.3 is used, which was the case for our local measurement. Therefore, we observe significant overhead from encryption. It is also worth noting that Quad9 and Google experienced higher latency at DoE (*i. e.*, DoH2 and DoT3) than Cloudflare and DNS4EU. This cannot be attributed to a suboptimal vantage point, as the latency of unencrypted DNS is similar to that of other resolvers. Also, the encryption algorithms do not cause such an additional latency [28], [29]. Therefore, it seems there are 1–2 RTTs (*i. e.*, $\approx$20–40 ms) of latency used for further tasks related to encryption or Starlink connections throttle after multiple round trips. We assume the first to be more likely.

Takeaway 3. DoE (*i. e.*, DoT3 and DoH2) in Starlink has a latency four to five times that of Do53. The measured variants of the DoE protocols require three round-trip, compared to one round-trip for Do53. Therefore, we observe a latency overhead of 1–2 round trips (*i. e.*, 20–40 ms) related to encryption.

IX. Conclusion

Our work visited various aspects of DNS performance in Starlink networks. To answer RQ 1 (*i. e.*, *default resolver choice in Starlink*), we found a resolver setting illustrated in Figure 2 that defaults to the resolver `34.145.127.1` that only resolves *starlink.com* and *spacex.com*. Upon subscription, users default to SpaceX-operated resolvers, which appear to offer the lowest latency among the studied resolvers. The major influences on latency (RQ 2) are primarily the length of the data path. Presumably, a resolver within a Starlink PoP offers the shortest possible data path. We also found that latencies in DNS are mostly determined by the user antenna to PoP latency (average $\approx$68.83 % of total latency). We do not assume server-side improvements to improve the overall DNS latency. The only way to improve DNS latency is to move the resolver closer to the user, which can be achieved by either improving Starlink's infrastructure or deploying satellite resolvers. However, the latter may prove difficult as many authoritative resolvers will remain in the terrestrial network. In such a case, recursive resolution still needs to traverse the satellite constellation, resulting in the sharp rise in the DNS latency observed in this work. If such a step is taken, we may observe that large providers can further improve latency, while smaller ones cannot.

Further, we wanted to study differences in public resolvers, as in RQ 3. We found that the studied resolvers (*e. g.*, Quad9 and DNS4EU) did not show significant differences. If differences were observed, they aligned with the resolvers' goals (*e. g.*, global service or focus on the European region). Further, the study examines the effects of different DNS protocols (RQ 4). Most protocols align with the number of round trips, except for DoE, which incurs roughly 1–2 additional round trips. This seems to be related to encryption. Further, we found that the used query types (A and AAAA) do not affect latency. Caching, however, does make a difference, but only to an expected extent.

Overall, we provided a comprehensive overview of existing DNS infrastructure from the perspective of Starlink users. Further studies should examine the actual effects of SpaceX-operated resolvers to better understand their architectural properties.

X. Ethics and Use of AI

AI was used to improve and correct grammar. Specifically, Anthropics Claude Opus 4.5 and 4.6, and Grammarly were used throughout the text body, and the text was double-checked for intended content. It aligns with IEEE Author Guidelines for Artificial Intelligence (AI)-Generated Text. The study does not raise ethical concerns.

References

[1] M. Zhou, X. Zhang, S. Hao, X. Yang, J. Zheng, G. Chen, and W. Dou, "Regional IP anycast: Deployments, performance, and potentials," in *Proceedings of the ACM SIGCOMM 2023 Conference, ACM SIGCOMM 2023, New York, NY, USA, 10-14 September 2023*, H. Schulzrinne, V. Misra, E. Kohler, and D. A. Maltz, Eds. ACM, 2023, pp. 917–931. [Online]. Available: https://doi.org/10.1145/3603269.3604846

[2] P. V. Mockapetris, "Domain names - implementation and specification," *RFC*, vol. 1035, pp. 1–55, 1987. [Online]. Available: https://doi.org/10.17487/RFC1035

[3] J. Dickinson, S. Dickinson, R. Bellis, A. Mankin, and D. Wessels, "DNS transport over TCP - implementation requirements," *RFC*, vol. 7766, pp. 1–19, 2016. [Online]. Available: https://doi.org/10.17487/RFC7766

[4] Z. Hu, L. Zhu, J. S. Heidemann, A. Mankin, D. Wessels, and P. E. Hoffman, "Specification for DNS over transport layer security (TLS)," *RFC*, vol. 7858, pp. 1–19, 2016. [Online]. Available: https://doi.org/10.17487/RFC7858

[5] P. E. Hoffman and P. McManus, "DNS queries over HTTPS (doh)," *RFC*, vol. 8484, pp. 1–21, 2018. [Online]. Available: https://doi.org/10.17487/RFC8484

[6] R. Bose, S. Fadaei, N. Mohan, M. M. Kassem, N. Sastry, and J. Ott, "It's a bird? it's a plane? it's cdn!: Investigating content delivery networks in the LEO satellite networks era," in *Proceedings of the 23rd ACM Workshop on Hot Topics in Networks, HOTNETS 2024, Irvine, CA, USA, November 18-19, 2024*. ACM, 2024, pp. 1–9. [Online]. Available: https://doi.org/10.1145/3696348.3696879

[7] R. Bose, N. Mohan, and J. Ott, "Poster: Twinkle, twinkle, streaming star: Illuminating CDN performance over starlink," in *Proceedings of the 2024 ACM on Internet Measurement Conference, IMC 2024, Madrid, Spain, November 4-6, 2024*, N. Vallina-Rodriguez, G. Suarez-Tangil, D. Levin, and C. Pelsser, Eds. Madrid, Spain: ACM, 2024, pp. 759–760. [Online]. Available: https://doi.org/10.1145/3646547.3689666

[8] R. Richter, V. Ververis, and V. Bajpai, "Breaking through the clouds: Performance insights into starlink's latency and packet loss," in *2025 IFIP Networking Conference (IFIP Networking 2025), Limassol, Cyprus, May 26-29, 2025*, 2025.

[9] F. Michel, M. Trevisan, D. Giordano, and O. Bonaventure, "A first look at starlink performance," in *Proceedings of the 22nd ACM Internet Measurement Conference, IMC 2022, Nice, France, October 25-27, 2022*, C. Barakat, C. Pelsser, T. A. Benson, and D. R. Choffnes, Eds. Nice, France: ACM, 2022, pp. 130–136. [Online]. Available: https://doi.org/10.1145/3517745.3561416

[10] Starlink, "What ip addresses does starlink assign?" Jan. 2026. [Online]. Available: https://starlink.com/support/article/1192f3ef-2a17-31d9-261a-a59d215629f4

[11] N. Mohan, A. E. Ferguson, H. Cech, R. Bose, P. R. Renatin, M. K. Marina, and J. Ott, "A multifaceted look at starlink performance," in *Proceedings of the ACM on Web Conference 2024, WWW 2024, Singapore, May 13-17, 2024*, T. Chua, C. Ngo, R. Kumar, H. W. Lauw, and R. K. Lee, Eds. ACM, 2024, pp. 2723–2734. [Online]. Available: https://doi.org/10.1145/3589334.3645328

[12] S. Sassalla, V. Ververis, and V. Bajpai, "A first look on discovery of designated resolvers," in *2025 IFIP Networking Conference (IFIP Networking 2025), Limassol, Cyprus, May 26-29, 2025*, 2025.

[13] Z. Tsiatsikas, G. Karopoulos, and G. Kambourakis, "Measuring the adoption of TLS encrypted client hello extension and its forebear in the wild," in *Computer Security. ESORICS 2022 International Workshops - CyberICPS 2022, SECPRE 2022, SPOSE 2022, CPS4CIP 2022, CDT&SECOMANE 2022, EIS 2022, and SecAssure 2022, Copenhagen, Denmark, September 26-30, 2022, Revised Selected Papers*, ser. Lecture Notes in Computer Science, S. K. Katsikas, F. Cuppens, C. Kalloniatis, J. Mylopoulos, F. Pallas, J. Pohle, M. A. Sasse, H. Abie, S. Ranise, L. Verderame, E. Cambiaso, J. M. Vidal, M. A. S. Monge, M. Albanese, B. Katt, S. Pirbhulal, and A. Shukla, Eds., vol. 13785. Springer, 2022, pp. 177–190. [Online]. Available: https://doi.org/10.1007/978-3-031-25460-4_10

[14] Y. Hauri, D. Bhattacherjee, M. Grossmann, and A. Singla, ""internet from space" without inter-satellite links," in *HotNets '20: The 19th ACM Workshop on Hot Topics in Networks, Virtual Event, USA, November 4-6, 2020*, B. Y. Zhao, H. Zheng, H. V. Madhyastha, and V. N. Padmanabhan, Eds. ACM, 2020, pp. 205–211. [Online]. Available: https://doi.org/10.1145/3422604.3425938

[15] S. Fadaei, S. Tiwari, A. Taneja, S. Bhushan, M. M. Kassem, A. Raman, D. Bhattacherjee, L. Qiu, A. Woodward, and N. Sastry, "Leoscope: Building a global testbed for low-earth orbit satellite networks," *Comput. Commun. Rev.*, vol. 55, no. 2, pp. 13–21, 2025. [Online]. Available: https://doi.org/10.1145/3750832.3750835

[16] J. Garcia, S. Sundberg, and A. Brunström, "TCP congestion control performance over starlink," in *Proceedings of the 2025 Applied Networking Research Workshop, ANRW 2025, Madrid, Spain, 22 July 2025*. ACM, 2025, pp. 70–77. [Online]. Available: https://doi.org/10.1145/3744200.3744760

[17] Z. Lai, Z. Li, Q. Wu, H. Li, J. Li, X. Xie, Y. Li, J. Liu, and J. Wu, "Leocc: Making internet congestion control robust to LEO satellite dynamics," in *Proceedings of the ACM SIGCOMM 2025 Conference, SIGCOMM 2025, São Francisco Convent, Coimbra, Portugal, September 8-11, 2025*, M. Curado, C. E. Rothenberg, G. Porter, and S. Kandula, Eds. ACM, 2025, pp. 129–146. [Online]. Available: https://doi.org/10.1145/3718958.3750491

[18] W. X. Zheng, A. Taneja, M. Masood, A. Sabnis, R. K. Sitaraman, and D. Vasisht, "Starcdn: Moving content delivery networks to space," in *Proceedings of the ACM SIGCOMM 2025 Conference, SIGCOMM 2025, São Francisco Convent, Coimbra, Portugal, September 8-11, 2025*, M. Curado, C. E. Rothenberg, G. Porter, and S. Kandula, Eds. ACM, 2025, pp. 948–962. [Online]. Available: https://doi.org/10.1145/3718958.3754345

[19] R. Bose, J. Zhao, T. Shreedhar, J. Pan, and N. Mohan, "Investigating web content delivery performance over starlink," *CoRR*, vol. abs/2510.13710, 2025. [Online]. Available: https://doi.org/10.48550/arXiv.2510.13710

[20] RIPE Atlas, "Starting your own measurements (user-defined measurements): Quotas," Jan. 2026. [Online]. Available: https://atlas.ripe.net/docs/getting-started/user-defined-measurements#quotas

[21] V. Le Pochat, T. Van Goethem, S. Tajalizadehkhoob, M. Korczyński, and W. Joosen, "Tranco: A research-oriented top sites ranking hardened against manipulation," in *Proceedings of the 26th Annual Network and Distributed System Security Symposium*, ser. NDSS 2019, Feb. 2019.

[22] G. C. M. Moura, J. S. Heidemann, R. de Oliveira Schmidt, and W. Hardaker, "Cache me if you can: Effects of DNS time-to-live," in *Proceedings of the Internet Measurement Conference, IMC 2019, Amsterdam, The Netherlands, October 21-23, 2019*. ACM, 2019, pp. 101–115. [Online]. Available: https://doi.org/10.1145/3355369.3355568

[23] root servers.org, "Root servers," Jan. 2026. [Online]. Available: https://root-servers.org

[24] T. V. Doan, I. Tsareva, and V. Bajpai, "Measuring DNS over TLS from the edge: Adoption, reliability, and response times," in *Passive and Active Measurement - 22nd International Conference, PAM 2021, Virtual Event, March 29 - April 1, 2021, Proceedings*, ser. Lecture Notes in Computer Science, O. Hohlfeld, A. Lutu, and D. Levin, Eds., vol. 12671. Springer, 2021, pp. 192–209. [Online]. Available: https://doi.org/10.1007/978-3-030-72582-2_12

[25] S. García, K. Hynek, D. Vekshin, T. Cejka, and A. Wasicek, "Large scale measurement on the adoption of encrypted DNS," *CoRR*, vol. abs/2107.04436, 2021. [Online]. Available: https://arxiv.org/abs/2107.04436

[26] C. Lu, B. Liu, Z. Li, S. Hao, H. Duan, M. Zhang, C. Leng, Y. Liu, Z. Zhang, and J. Wu, "An end-to-end, large-scale measurement of dns-over-encryption: How far have we come?" in *Proceedings of the Internet Measurement Conference, IMC 2019, Amsterdam, The Netherlands, October 21-23, 2019*. ACM, 2019, pp. 22–35. [Online]. Available: https://doi.org/10.1145/3355369.3355580

[27] S. Ma, Y. C. Chou, H. Zhao, L. Chen, X. Ma, and J. Liu, "Network characteristics of LEO satellite constellations: A starlink-based measurement from end users," in *IEEE INFOCOM 2023 - IEEE Conference on Computer Communications, New York City, NY, USA, May 17-20, 2023*. IEEE, 2023, pp. 1–10. [Online]. Available: https://doi.org/10.1109/INFOCOM53939.2023.10228912

[28] C. Silva, V. A. Cunha, J. P. Barraca, and R. L. Aguiar, "Analysis of the cryptographic algorithms in iot communications," *Inf. Syst. Frontiers*, vol. 26, no. 4, pp. 1243–1260, 2024. [Online]. Available: https://doi.org/10.1007/s10796-023-10383-9

[29] E. Kwesigabo, "Evaluating the impact of latency in a virtual private network performance using different encryption algorithm and hashing," *The Journal of Informatics*, vol. 4, no. 1, Dec. 2024. [Online]. Available: https://journals.iaa.ac.tz/index.php/tji/article/view/245

[30] satellitemap.space, "Live satellite map," Jan. 2026. [Online]. Available: https://satellitemap.space

[31] vasilis ververis, S. Sassala, F. Roth, and V. Bajpai, "Path to encrypted DNS with DDR: adoption, configuration patterns, and privacy implications," *Proc. Priv. Enhancing Technol.*, vol. 2025, no. 4, pp. 465–484, 2025. [Online]. Available: https://doi.org/10.56553/popets-2025-0140

[32] B. Wang, X. Zhang, S. Wang, L. Chen, J. Zhao, J. Pan, D. Li, and Y. Jiang, "A large-scale ipv6-based measurement of the starlink network," *CoRR*, vol. abs/2412.18243, 2024. [Online]. Available: https://doi.org/10.48550/arXiv.2412.18243

[33] W. Ali, C. Fang, and A. Khan, "A survey on the state-of-the-art CDN architectures and future directions," *J. Netw. Comput. Appl.*, vol. 236, p. 104106, 2025. [Online]. Available: https://doi.org/10.1016/j.jnca.2025.104106

[34] J. Garcia, M. Beckerle, S. Sundberg, and A. Brunström, "Modeling and predicting starlink throughput with fine-grained burst characterization," *Comput. Commun.*, vol. 234, p. 108090, 2025. [Online]. Available: https://doi.org/10.1016/j.comcom.2025.108090