A Long-term View of DNS over QUIC Adoption and its Performance Impact on YouTube Streaming

Jayasree Sengupta*, Mike Kosek†, Justus Fries†, Veronika Kitsul‡, and Vaibhav Bajpai§
*Indian Institute of Information Technology, Allahabad, India [jayasree@iiita.ac.in]
†Technical University of Munich, Germany [{kosek | fries}@in.tum.de]
‡Princeton University, USA [vk6976@princeton.edu]
§Hasso Plattner Institute & University of Potsdam, Germany [vaibhav.bajpai@hpi.de]

***Abstract*—YouTube contributes the largest share of global video traffic on the Internet, making it an important use case for understanding the impact of evolving DNS protocol choices on video streaming performance. Although traditional DNS over UDP (DoUDP) offers low latency, it lacks modern transport features. Encrypted DNS protocols such as DNS over TLS (DoT) and DNS over HTTPS (DoH) improve protocol robustness but suffer from higher latency due to their underlying transport and encryption protocols with multi-RTT handshakes. However, recently standardized DNS over QUIC (DoQ) aims to combine the best of both worlds by leveraging the transport efficiency of QUIC while ensuring DNS privacy. In this paper, we present the first comprehensive long-term measurement study of DoQ adoption and evaluate its performance implications for YouTube video streaming. We collect data through weekly scans of the IPv4 address space over a two-year period to assess the adoption of the protocol. Our results show that DoQ adoption by public DNS resolvers has steadily increased and plateaued over 25 months. Using seven globally distributed vantage points, our *video performance* measurements shows that DoQ's DNS lookup time increases by only 1.5% in the median while video startup delay increases by less than 1% compared to DoUDP. In particular, in about 40% of the cases, DoQ yields faster video startup times than DoUDP. These findings position DoQ as a technically efficient DNS protocol, well suited for modern, high-demand performance-sensitive applications such as video streaming.**

***Index Terms*—QUIC, Encrypted DNS, Internet measurement, YouTube video streaming.**

I. Introduction

The Domain Name System (DNS) plays a foundational role in nearly all Internet-based activities, including Web browsing and video streaming. Traditional unencrypted DNS protocols like DNS over UDP (DoUDP) and DNS over TCP (DoTCP) remain widely used due to their low latency. However, they lack modern transport-layer capabilities such as built-in encryption, multiplexing support, and congestion control, making DNS requests and response vulnerable to eavesdropping and on-path manipulations [1]. Hence, a user profile can be created and tracked with only having access to the user's DNS traffic [2], [3], [4]. Such data can then be leveraged to provide personalized recommendations, targeted advertising [5] and can also be exploited by attackers, leading to major privacy breaches. To address these limitations, encrypted DNS protocols—including DNS over TLS (DoT) and DNS over HTTPS (DoH) have been introduced, offering enhanced security and protocol robustness. As these protocols have been extensively studied (see II) in terms of response times [6], [7], [8], [9], [10] and impact on Web performance [11], [8], [12], it has become clear that both DoT and DoH are constrained by the round trips required for the handshakes of the underlying transport (TCP) and encryption (TLS) protocols. These challenges are addressed by the QUIC transport protocol [13], [14], [15], which overcomes handshake limitations by combining the transport and encryption handshake into a single round trip. Consequently, DNS over QUIC (DoQ) [16] evolved, which improves over both DoH and DoT in terms of latency.

Currently, YouTube video streaming generates one of the largest sources of Internet traffic in 2025 [17]. According to a latest report [18], YouTube has more than 2.6 billion active users with over 122 million people visiting YouTube everyday. Another report [19] published in 2022 indicates that 1 billion hours of YouTube video are streamed per day. As such, YouTube is a lucrative choice for profiling users across the world, especially for monitoring video streaming activities. YouTube uses Dynamic Adaptive Streaming over HTTP (DASH) [20] to stream videos which means chunks of video are fetched through HTTPS requests. Since every HTTPS request relies on DNS to resolve domain names before redirecting the client to a geographically proximate Content Delivery Network (CDN) replica for fetching video chunks, video streaming consequently tightly couples DNS with HTTPS [8]. In order to leverage the benefits offered by the modern encrypted DNS protocols during DNS resolution, we replace DoUDP with varieties of encrypted DNS (DoE): DoT [21], DoH [22] and DoQ [16] for YouTube streaming as shown in Fig. 1.

Given the scale of YouTube and its reliance on low-latency, real-time content delivery, even minor inefficiencies in DNS resolution can translate to noticeable user experience degradation. Hence, performance metrics such as lookup time, video startup delay, etc. are closely tied to delays incurred during DNS resolution. Thus, it is important to quantify the trade-offs between various performance metrics, for assessing its suitability in real world latency-sensitive applications. A comprehensive evaluation across all DNS variants provides a more holistic view of their suitability for high-demand use cases like video streaming and informs protocol selection decisions for ISPs, CDN operators, and browser vendors.

This paper builds on our earlier works [23], [24]. Contrary to the snapshot measurement in [23], here we collect DoQ

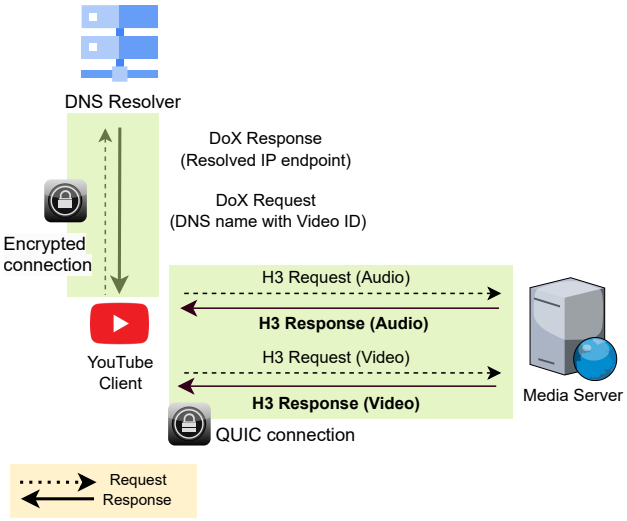

Fig. 1: Proposed mechanism of YouTube video streaming where DNS resolution happens over encrypted DNS, instead of unencrypted plain text (DoUDP).

data every week continuously over the course of 25 months from July, 2021 until July, 2023, to investigate its longitudinal adoption (see IV), reflecting the ongoing development and standardization process, during which DoQ implementations and services underwent rapid changes. Our study shows that DoQ has been widely accepted by the community, where, *AdGuard* [25] and *nextDNS* [26] already use DoQ in production systems for their DNS-based ad as well as tracker blocking services, offering publicly reachable DoQ servers and client implementations [27], [28]. While our previous study [24], evaluated DoQ and its impact on Web browsing across six vantage points, in this paper we conduct *video performance* measurements (see V) over a week across all vantage points to evaluate the performance of DoQ during YouTube video streaming and compare it with the corresponding DNS protocols: DoUDP, DoTCP, DoT, and DoH. As Web and YouTube are two different workloads, therefore, performance benefits seen in earlier work [24] for websites do not directly apply to other workloads (e.g. video streaming). Our measurement setup also modifies the DNS Proxy and Chromium (see: § III). The main contributions of this work are the measurement method built to perform the study and the findings from such measurements as detailed below:

- **Measurement Method –** We built a measurement method (see: § III) to study the adoption of DoQ since its inception and evaluate its performance while streaming YouTube videos. We first measure the adoption of DoQ by observing 8 DoQ/QUIC version pairs every fourth week over the course of 107 weeks. We then extend the study to analyze varieties of DoX protocols with H/3 when streaming Youtube videos. We measure the protocols across a diverse set of seven vantage points and a target list of 312 resolvers. The setup uses Google Chrome in headless mode with HTTP caching disabled, H/3 with enforced QUIC and session ticket caching to support session resumption. A fork of DNS proxy is used as the local stub resolver while the proxy is set as the local nameserver. The setup performs extensive DNS performance metrics logging to support analytics during connection reuse.

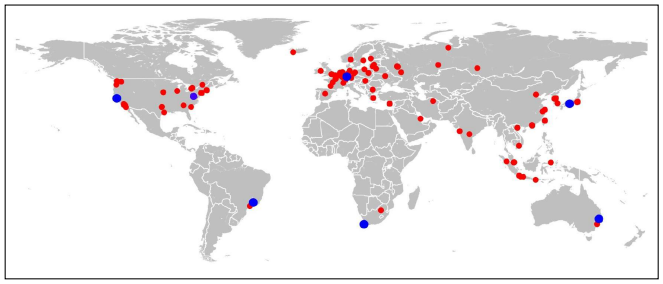

Fig. 2: Geographical distribution of the 312 DoX-verified resolvers (red dots) and seven vantage points (blue dots).

- **Findings –** We find that the adoption of DoQ by public DNS resolvers slowly increased until June 2022. In July 2022, the adoption almost halved, but since then has maintained an almost stable state (see: § IV). For the *video performance* measurements (see V), we show that DNS lookup time using DoQ is only minimally higher in comparison to DoUDP with $\approx$1.5% in the median. The median startup delay increase with DoQ over DoUDP is $< 1\%$ which continues up to the 90^{th} percentile. DoH on the other hand tends to perform worse under deteriorating network conditions with results varying from $3.9\% - 10.2\%$ from the median to the 90^{th} percentile. Normalizing the different protocols per vantage point, we find that for almost 40% of our samples, DoQ even performs better than DoUDP in terms of startup delay, thereby amortising the cost of encrypted DNS with DoQ for YouTube streaming.

We finally end with the concluding remarks in § VI.

II. Related Work

A. Encrypted DNS Protocols

While both DoT and DoH address the key issues of adding privacy to DNS [9], [1], [6], [29], they are inherently constrained by head-of-line-blocking and missing multiplexing support on the transport layer, as well as an additional connection establishment in comparison to UDP. These challenges were addressed with the standardization of QUIC [13], [14], [15] in early 2021. QUIC is a connection-oriented encrypted transport protocol, built on top of TLS 1.3 by using UDP as a substrate. QUIC features mandatory encryption, provides stream multiplexing, and mitigates transport-layer head-of-line blocking by enabling independent delivery across streams. It improves connection establishment time by combining the transport and encryption handshakes into a single round trip. Further, with QUIC 0-RTT, clients can send application data in the very first round trip of the connection, without requiring any other handshake to be completed beforehand. QUIC also provides reliable communication via flow control, congestion control, and loss detection mechanisms. Thus, DNS over QUIC (DoQ) [16], is a third attempt to improve DNS privacy with minimum latency by leveraging QUIC as the underlying protocol. In other words, although DoQ carries privacy properties similar to DoT and DoH, the latency characteristics

of DoQ is more similar to the unencrypted DoUDP [24], [30], [31]. With this objective, DoQ aims to obsolete all other currently used DNS protocols, which lack privacy and/or require more round-trips for handshakes—therefore promising to make DoQ the predominant successor.

Lyu *et al.* [32] survey the DNS encryption standards and literature between 2016 to 2021 looking at their adoption status, performance, benefits, and security issues. Several studies [9], [33] also highlight the evolution and adoption of encrypted DNS protocols over time. In [29], [23], the authors explicitly study the early stage adoption of DoQ, showing a steady increase with a considerable portion of the measurements indicating higher than expected handshake times; yet, DoQ still outperforms DoT and DoH. Further, a study by Kosek *et al.* [24] evaluates DoQ and its impact on Web browsing across multiple vantage points. More recently, Bielefeld *et al.* [34] presented a large-scale empirical analysis of DoQ vs. DNS-over-HTTP/3 (DoH/3) and characterize their support for different features such as session resumption and 0-RTT. From the perspective of privacy-preservation, Siby *et al.* [35] examine whether encrypted DNS traffic can protect users from traffic analysis-based monitoring [36] and censoring. The results indicate that DNS-based censorship is feasible even on encrypted DNS traffic. Contrarily, the study [37] aims at comparing the level of privacy leakage in encrypted DNS, specifically with respect to DoQ traffic. Lastly, [30], [31] evaluates the benefit of using QUIC to coalesce name resolution via DNS over QUIC (DoQ), and Web content delivery via HTTP/3 (H3) with 0-RTT. While encrypted DNS protocols such as DoQ provide confidentiality guarantees by design, as established in prior work, our study focuses on evaluating their performance implications in real-world applications rather than directly measuring privacy properties.

B. Video Streaming Measurements

A separate branch of literature (see: Table I) highlights secure and privacy-preserving content access for different video streaming platforms. In one such work, Rajan *et al.* [38] propose a Hierarchical Inner Product Encryption (HIPE) based system that offers multiple levels of data access control to end users based on their roles. It allows them to stream only those videos for which they have access. On the other hand, Hooman *et al.* [39] demonstrate privacy issues with Over-the-Top (OTT) services when using devices such as Amazon Fire Stick and Roku. They developed a system which automatically downloads OTT apps (e.g. YouTube), to interact with the devices while intercepting the network traffic and performing best-effort TLS interception. Feng *et al.* [40] present Silhouette, a real-time, lightweight video classification method suitable for ISP middle-boxes. It uses flow statistics for identifying video flows even when they are encrypted. Recent work [41] has also focused on the measurement and characterization of encrypted video streaming platforms, including the identification of different user platforms and traffic behavior in broadband networks.

Beyond encrypted DNS, a broader body of work has examined DNS from the perspective of network management, security, and operational visibility. For example, a recent study [42] has explored enterprise DNS asset mapping and monitoring using passive traffic analysis. In contrast to these works, our study focuses on the intersection of DNS transport evolution and application-level performance. Specifically, we combine a longitudinal analysis of DoQ adoption with an end-to-end evaluation of its impact on video streaming QoS and QoE. This cross-layer perspective distinguishes our work from prior DNS measurement and video streaming studies.

TABLE I: Existing Research on Privacy Enhancement of Encrypted DNS and Video Streaming Services

Topics	Research Focus	References
DNS Privacy	Survey on DNS Encryption Standards	[32]
Quantifying Privacy	Evolution and adoption of Encrypted DNS	[9], [33]
	Early stage adoption of DoQ	[29], [23]
Web Privacy	Impact of DoQ and DoH/3 on Web performance	[24], [34]
	QUIC connection coalescing using DoQ and H3	[30], [31]
Traffic Analysis	Privacy leakage in Encrypted DNS traffic	[35], [37]
Privacy-preserving Video Streaming	Role-based access control for video streaming	[38]
	Privacy issues in OTT services	[39]
	Encrypted video stream classification method	[40]
	Characterizing Video Streaming Platforms	[41]

III. Methodology

Our methodology consists of three main stages: (i) resolver discovery via Internet-wide scanning, (ii) protocol validation and filtering, and (iii) video performance measurements.

A. Discovery of Target Resolvers

1) Scanning: To assess and identify the adoption of DNS over QUIC (DoQ) resolvers worldwide, we issue weekly scans of the IPv4 address space starting from 2021-W27 (July 05–11). The scans are conducted from a single vantage point located in the research network of TUM, Germany targeting all proposed DoQ ports (`UDP 784`, `853`, and `8853` [16]). For this, we leverage the *ZMap* [43] network scanner. For comparison, we additionally target DoUDP port `UDP/53`, which we identify by leveraging the *ZMap*'s built-in DNS probing packet that queries an `A` record for `www.google.com` [44]. Since *ZMap* does not provide means for the identification of QUIC or DoQ, we issue a custom packet [45] that carries the `Initial` QUIC handshake frame with an invalid version number of `0` [46]: In this way, if the target operates a QUIC stack on the probed port, a `Version Negotiation` packet

This article has been accepted for publication in IEEE Transactions on Network and Service Management. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TNSM.2026.3688441

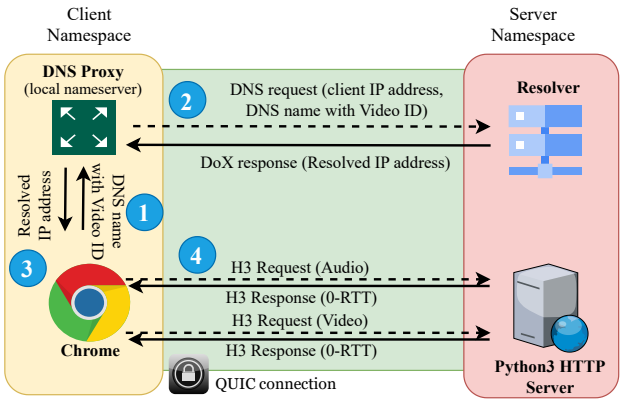

Fig. 3: Measurement setup used to evaluate performance of varieties of DoX with H/3 when streaming YouTube videos.

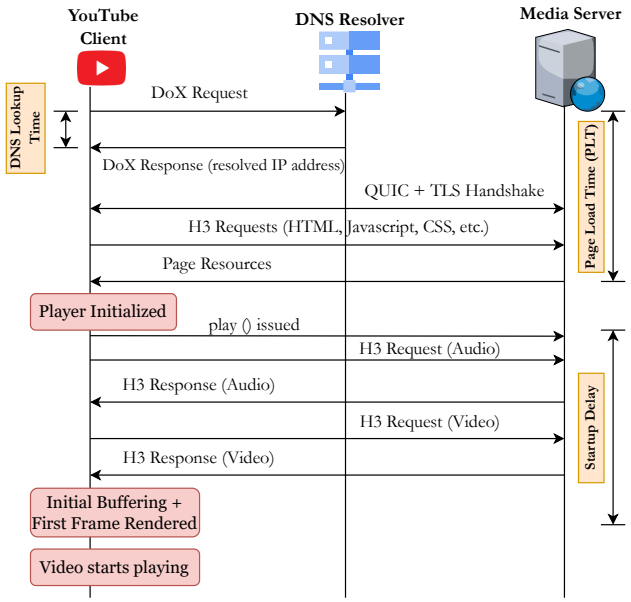

Fig. 4: A sequence diagram illustrating the YouTube video streaming workflow and where the different metrics are collected.

is triggered in response, which allows us to identify the target IP address as *QUIC-capable* on the respective port. This also prevents exhausting resources of the target [13] as such a packet does not produce state. However, note that other QUIC services, which are not necessarily DoQ, could be offered on the probed ports.

2) Target Validation: To further validate targets identified as *QUIC-capable* by the *ZMap* scans, we check if they actually support DNS over QUIC [47]. To do so, we offer the `doq Application-Layer Protocol Negotiation (ALPN)` identifiers (as required by the DoQ I-Ds [16]), which results in a list of *DoQ-capable* targets. As a final step, a connection to every *DoQ-capable* target on all proposed DoQ ports `UDP/784`, `UDP/853`, as well as `UDP/8853`, is established [48]: For these connections, we offered the QUIC version `draft-34` in our `Initial` frame until 2021-W42, while support for version `1` was added in 2021-W43. Overall, our client supports the QUIC versions `draft-34`, `-32` and `-29` since the start of our study, as well as version `1` later on; hence, the client can respond to `Version Negotiation` packets if issued by the resolvers. For DoQ, we offer versions in the order of final RFC [16] and `draft-11` to `draft-00` [16], for which we added support for new versions within 2 weeks of the `draft` or final RFC release. By issuing the highest QUIC and DoQ protocol versions supported by our client first, we ensure that we negotiate the highest shared protocol versions between our client and the target resolver. With this, we record the negotiated QUIC and DoQ versions, as well as the X.509 certificate offered by each *DoQ-capable* target, creating the final list of *DoQ-verified* resolvers.

B. Protocol Validation and Filtering

Using the methodology described above, we identify 1,216 DoQ resolvers. To enable a comparison of DoQ to other DNS protocols, we evaluate their support of DoUDP, DoTCP, DoT, and DoH. For this purpose, we use *DNSPerf*, an open-source DNS measurement tool supporting all stated protocols [49] to optimistically query the resolvers. Of the 1,216 identified DoQ resolvers, we find that 548 support DoUDP, 706 DoTCP, 1,149 DoT, and 732 DoH while their full intersection results in 312 verified DoX-verified resolvers (shown in Fig. 2), which support all five target DNS protocols (DoQ, DoH, DoT, DoTCP, DoUDP) as established by our preliminary study [24]. While we acknowledge that public DNS resolvers often leverage IP anycast, we cross-reference anycast IP addresses used in related work [6], [9], [10], [11], [12], [50], although without finding an overlap. Using an IPv4 geolocation lookup service [51], we find that the majority of these resolvers are located in Europe ($\approx 41.7\%$) and Asia ($\approx 41.0\%$), followed by North America ($\approx 15.4\%$), while Africa, Oceania, and South America each account for $\approx 0.6\%$ of the resolver set.

C. Video Measurements

1) Vantage Points: To assess the impact of DoQ in comparison to DoUDP, DoTCP, DoT, and DoH on video performance (YouTube video streaming), we perform distributed measurements using 7 *Amazon EC2* instances while targeting these 312 DoX-verified resolvers as represented in Fig. 2. The seven virtual machines are located in North America (California and Virginia, USA), South America (São Paulo, Brazil), Europe (Frankfurt, Germany), Africa (Cape Town, South Africa), East Asia (Osaka, Japan) and Oceania (Sydney, Australia). The measurements were run once a day starting at 00:00 UTC from 18th-24th April (covering seven days). Following a similar strategy as our earlier work [24], we compare DNS transports relative to each-other within the same vantage point–resolver pair. By evaluating all DNS transports for the same vantage point–resolver pairs, we reduce variability due to endpoint differences; however, we do not explicitly capture path-level dynamics such as routing changes or congestion. Although the geographic distance between a vantage point and a resolver influences the absolute lookup latency, our analysis emphasizes relative differences across transports for a fixed VP–resolver combination. Since all protocols are tested between the same endpoints, they traverse the same underlying network path and are subject to identical propagation delay, routing characteristics, and congestion conditions. Consequently, path-related

factors affect all transports in a similar way and do not distort the cross-protocol comparisons. This follows standard practice in DNS transport measurement studies, including our recent study.

2) Settings: We intentionally fix hardware, browser configuration, and local network settings across all measurement vantage points to isolate the performance contribution of the DNS transport protocols. Varying device types, CPU capabilities, or artificially throttled network conditions would introduce confounding effects and make cross-protocol comparison non-repeatable. Instead, heterogeneity arises naturally through our globally distributed vantage points, each operating under distinct routing paths, latencies, and peering characteristics. This allows us to study protocol behavior in diverse real-world conditions while maintaining experimental control. Our experiments interact with YouTube's real-world delivery system, including Dynamic Adaptive Streaming over HTTP (DASH) and redirection to the closest CDN replicas. By using the real Internet, our study accounts for real-world phenomena like uncached domain name requests and varying throughput between VPs and media servers.

3) Measurement Setup: Our methodology represents the usage of YouTube via direct links where on clicking, the DNS resolution is performed as explained below, and the video starts playing. Our measurement setup (see: Fig. 3) considers two different YouTube videos [52], [53] in order to compare the domain names used for fetching video chunks. The videos are always measured for five seconds with the video quality set to both *720p* and *480p*. These two videos are specifically chosen since they are ad-free throughout the testing. Ads introduce additional DNS and HTTP requests which complicate controlled comparisons across DNS transports. Using two globally ad-free videos ensures experimental repeatability across vantage points. As such, adding more videos would likely yield redundant data rather than new insights. Instead, our study brings diversity by adding seven global vantage points measuring 312 DoX verified resolvers leading to $>$11k successful DoQ measurements alone. The measurement script first iterates over the 312 DNS resolvers to confirm whether they are reachable, and then iterates over two YouTube videos. To effortlessly switch between the measured protocols (DoT, DoH, DoQ, DoTCP and lastly DoUDP) without needing the browser to support the specific protocol, a fork of DNS proxy is used as the local stub resolver. For each protocol, DNS proxy is run with the server set as the upstream resolver. DNS proxy listens for queries on a localhost address on port 53 and proxies them to the upstream resolver server. To force the browser to use DNS proxy as its stub resolver, the proxy is set as the local nameserver. The UDP source port is anchored to guarantee that a post-reset QUIC connection to the same server will have a handshake time of one round-trip. Similar to the measurement setup built in prior works [24], we incorporate support for session resumption by adding a session ticket cache. This ensures that after a reset, the next QUIC handshake to the DNS server does not require the server to resend its certificate again. For each *[server, video, protocol]* combination, the Chrome browser is opened using *Selenium* and *Chromedriver* to capture the different performance metrics. Chrome is set to run in headless mode with HTTP caching disabled while HTTP/3 (H3) with QUIC is enforced to ensure every video chunk is requested over H/3. The target website of the measurement is an *iframe* served by a HTTP server on localhost over HTTP without encryption. *iframe* API is an established method of mimicking the playback of a single video as already in use by previous studies [54], [55]. Our *iframe* code is based on the reference implementation [56] by YouTube. The *iframe* API allows setting the video ID to start playback programmatically, DNS performance metrics logging is also implemented to support proper recording of exchange duration.

4) Metrics Used: In our evaluation, we consider three metrics that capture different aspects of DNS and video performance as follows: (a) DNS lookup time, which includes query latency and protocol handshakes, (b) Page Load Time (PLT) which reflects all DNS and HTTP requests performed before playback begins, and (c) Startup delay which measures the time from initiating playback until the first frame is rendered. Fig. 4 illustrates the YouTube streaming workflow and highlights where each metric is measured. These metrics jointly allow us to evaluate both transport-efficiency effects and user-perceived performance. There are two points at which these metrics are recorded: using the Javascript embedded in the website serving the *iframe* and using Javascript executions in *Selenium*. Finally, the measurement runs where the *google-video* subdomain changes within a run are removed to reduce the effect of uncached domain name requests.

The optimal quality determined by a YouTube player inside an *iframe* can be influenced by setting the size of the player. While an *iframe* loads, multiple HTTP requests are issued to various hosts. The page load time (PLT) covers all of the DNS queries and HTTP requests related to these hosts and specifically excludes any video playback. The lookup time consists of the query time (i.e. time from sending the request until receiving the response) and potentially a handshake of the transport protocol. *Initial quality* is defined as the quality change event that is logged by the *iframe* API right before the playback start event is captured. Finally, *startup delay* refers to the time lapse between programmatically starting the playback to the player actually starting to play the video after initial buffering.

Ethical Considerations: To adhere to ethical principles and minimize the impact of our active scans, we follow best practices of the Internet measurement community [57], [58], [59]: For the longitudinal analysis, we restrict our Internet-wide scans to verify DoX resolvers once every week over a specific vantage point in order to limit outbound traffic. To allow targets to opt-out of our measurements, we display contact information and a description about the intent of our measurements on a webpage reachable via the IP address of each vantage point. Further, we only target publicly reachable IP addresses and honor opt-out requests from previous studies by maintaining a University-wide shared blocklist with the excluded targets.

Reproducibility: In order to enable the reproduction of our findings [60], we have made the developed tools, the raw data of our measurements, and the analysis scripts for the work

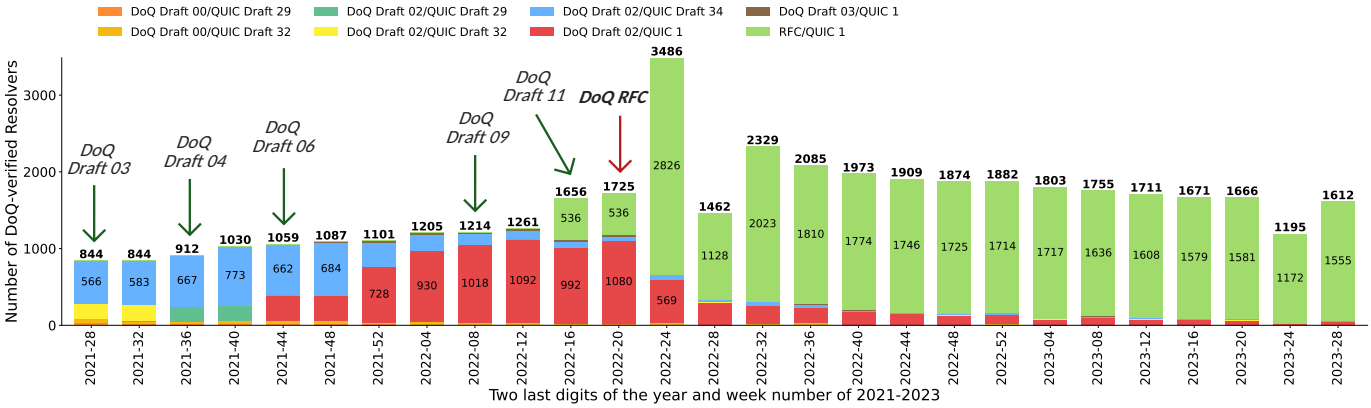

Fig. 5: Number of *DoQ-verified* resolvers per week number from July, 2021 to July, 2023 grouped by negotiated DoQ and QUIC version. Support for QUIC version `1` was added in 2021-W43. Also, highlights release of different DoQ drafts and its standardization timeline.

publicly available[1].

IV. Long-Term View on Adoption of DoQ

To study the adoption of DoQ on resolvers worldwide, we issue weekly scans of the IPv4 address space over the course of 25 months, i.e. 107 weeks, as detailed in § III. In our scans, we record the negotiated QUIC and DoQ versions, as well as the X.509 certificates offered by the target resolvers that support DoQ, for which we also determine the announcing Autonomous Systems (ASes) and geolocations. Overall, we find 1,851 unique X.509 certificates over this span.

Adoption of QUIC and DoQ Versions: In our scans and in the verification process, we target all proposed DoQ ports `UDP/784`, `UDP/853`, and `UDP/8853`. The DoQ drafts `-00` and `-01` state that port `UDP/784` *MAY* be used for experimentation. `draft-02` defined `UDP/8853` for usage as experimentation as well as for reservation at the Internet Assigned Numbers Authority (IANA). This was changed in `draft-03`, where port `UDP/784` was again stated for experimentation usage; ultimately, `UDP/853` has been established as the final port for reservation at IANA [16].

Fig. 5 presents the *DoQ-verified* resolvers per week, grouped by negotiated DoQ and QUIC version. Overall, we observe that the number of *DoQ-verified* resolvers rises steadily until 2022-W24: Starting with 844 resolvers in 2021-W28 (July 05–11), we see an increase to 3,486 verified resolvers in 2022-W24 (June 13–19). The number of *DoQ-verified* resolvers then sees a drastic drop in 2022-W28, but eventually maintains an almost steady number with 1,612 verified resolvers in 2023-W28 (July 08–14). This pronounced increase in the number of verified DoQ resolvers shortly after its standardization in May 2022 can be attributed to operators temporarily enabling DoQ for validation and interoperability testing following the publication of the final RFC. This interpretation is also supported by the large number of newly observed DoQ resolvers shown in Fig 6 during 2022-W24. Moreover, AdGuard Home (AGH), a widely used open-source DNS resolver, changed its default QUIC/DoQ version pairing around this time, leading to many AGH-based deployments suddenly becoming reachable via DoQ. However, we hypothesize that the sudden sharp decline (nearly halfing) around July 2022 reflects operational churn among early adopters. Many initial DoQ deployments were based on draft versions (e.g., doq-i02/03) that became incompatible with the standardized protocol, necessitating updates or temporary deactivation. In addition, AGH reverted or adjusted several DoQ-related configuration defaults during this period, which would have reduced the number of resolvers successfully negotiating DoQ with our client. Finally, the consistent upward trend, followed by saturation can be interpreted as a shift from experimental deployments toward stable, production-grade adoption. Larger resolver operators increasingly enabled standardized DoQ support, while smaller or prototype deployments disappeared. Together, these observations suggest that the DoQ ecosystem has progressed from early-stage experimentation to a more mature and stable operational phase. This observation is further supported by concurrent shifts in protocol version adoption and alignment with the standardization timeline (RFC 9250), as illustrated by the vertical markers in Figs. 5 and 6.

To analyze it further, we have measured the churn in the *DoQ-verified* resolvers (see: Fig. 6) over the entire duration of the measurement study. We observe that right after the standardization of DoQ in May, 2022, $\approx 57\%$ new resolvers joined in 2022-W24 (see: Fig. 6a), after which we again see a drastic drop. We speculate that, right after the standardization, several operators turned on DoQ to experiment its operation in the real world, resulting in such a quick spike that later disappeared. However, we must highlight that a steady number ($>$1000, median: 1043) of resolvers continue to support DoQ (cyan bars, Fig. 6a) demonstrating that the operations community continues offering DoQ in production. Also, a fairly consistent number of new resolvers ($\approx$ 200, median: 242) supporting DoQ (light green bar, Fig. 6a) were added every week throughout the entire duration of the study, which reflects the growing popularity of DoQ implementation amongst operators. After we added support for QUIC version 1 [13] in 2021-W43, we observe a steady usage of `DoQ Draft 02/QUIC 1` (red bars) until 2021-W50, followed by a steep increase until 2022-W12. We again observe a steady trend till 2022-W20, followed by a steep decrease

[1]https://github.com/Sree2021/TNSM-2026-YouTube

This article has been accepted for publication in IEEE Transactions on Network and Service Management. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TNSM.2026.3688441

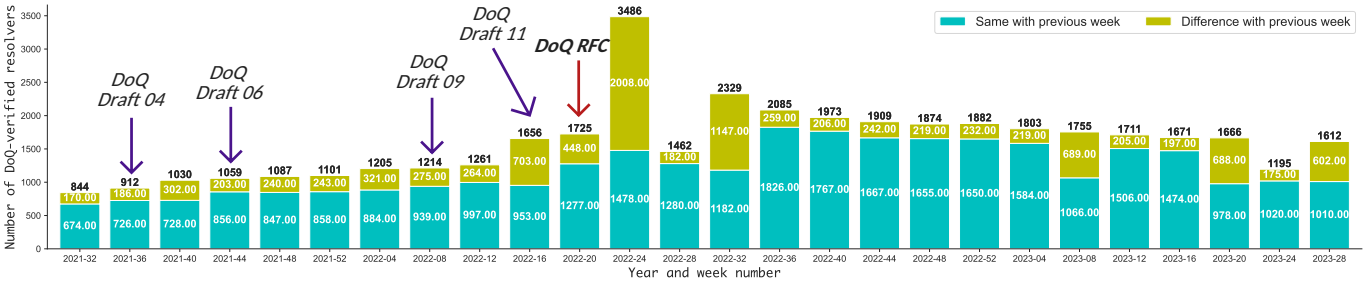

(a) With respect to previous week.

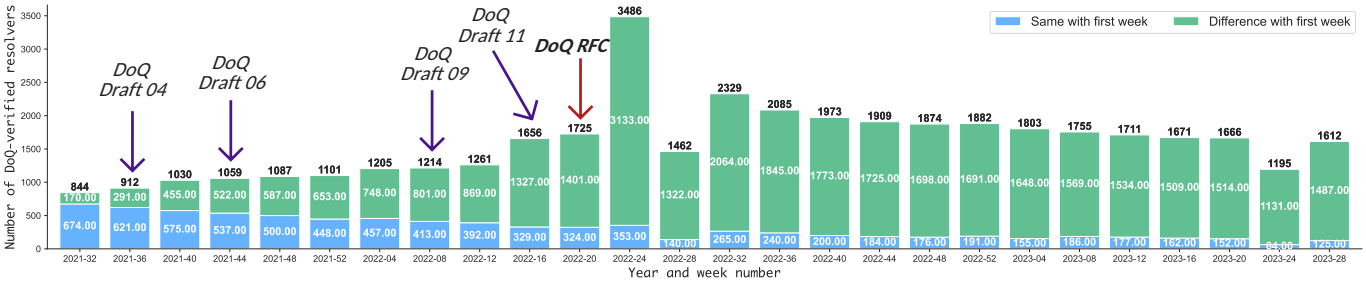

(b) With respect to first week (i.e. 2021-W28).

Fig. 6: Churn of *DoQ-verified* resolvers per week from August, 2021 to July, 2023. Vertical lines indicate key DoQ draft releases and RFC 9250.

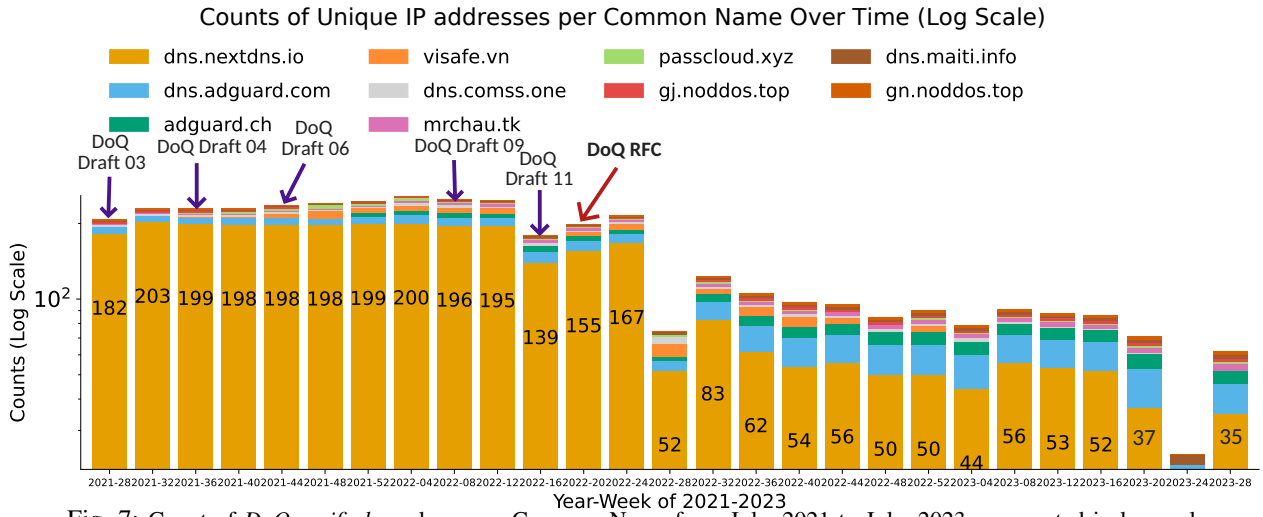

Fig. 7: Count of *DoQ-verified* resolvers per Common Name from July, 2021 to July, 2023 represented in log scale.

until 2023-W08. Similarly, after adding support for DoQ RFC shortly before the RFC release, we observe a steady usage of `RFC/QUIC 1` (light green bars) until 2022-W20, followed by a sudden steep rise in 2022-W24 (June 10–16), i.e. right after the standardization of DoQ [16]. There is a steep descend in 2022–W28, followed by an almost steady decrease until 2023-W28. Analyzing this observation, we find that the open source DNS server implementation *AdGuard Home (AGH)* [61] changed the default DoQ/QUIC pair from `DoQ Draft 02/QUIC Draft 34` (blue bars) to `DoQ Draft 02/QUIC 1` (red bars) starting 2021-W51 [62], matching the pattern we observe. In addition, we find indications of the usage of *AGH* by the updated resolvers within the *Common Names* of their X.509 certificates, and also identify multiple of the updated resolvers to be running *AGH* through random sampling. We also observe (Fig. 7) the count of unique IP addresses per common name suddenly dropped in 2022-W28 and then steadily decreased over time, matching the patterns observed in Fig. 5. Hence, we attribute the observed changes in usage of DoQ between 2021-W51 and 2023-W28 to these facts.

Although we offer a total combination of 52 DoQ/QUIC version pairs as of 2023-W28 (see § III), we observe only 8 pairs across all measurements, with the majority being `DoQ RFC/QUIC 1` (light green bars, 3181 (82.2%)) in 2022-W24. Also, when we look at the churn in the *DoQ-verified* resolvers with respect to the first week of our study, i.e. 2021-W28 (see: Fig. 6b), we see that only $\approx 9\%$ of the resolvers are

TABLE II: Number of *DoQ-verified* resolvers in censored/surveillance countries, distributed by network type as of July 10-16, 2023. China has the most number of *DoQ-verified* resolvers both within content delivery networks and ISPs.

Country	# Resolvers	Network Type			
		Content	ISP	Others	Unknown
China	127	91 (71.7%)	21 (16.5%)	0	15 (11.8%)
South Korea	51	35 (68.6%)	12 (23.5%)	1 (2%)	3 (5.9%)
Russia	50	9 (18%)	16 (32%)	2 (4%)	23 (46%)
Australia	14	13 (92.9%)	1 (7.1%)	0	0
Vietnam	12	0	6 (50%)	0	6 (50%)
Iran	3	1 (33.3%)	0	0	2 (66.7%)
Malaysia	3	3 (100%)	0	0	0
Norway	2	0	2 (100%)	0	0
Bangladesh	2	0	1 (50%)	0	1 (50%)
United Arab Emirates	1	1 (100%)	0	0	0
Kazakhstan	1	1 (100%)	0	0	0
Belarus	1	0	0	1 (100%)	0
TOTAL	**267**	**154 (57.7%)**	**59 (22.1%)**	**4 (1.5%)**	**50 (18.7%)**

TABLE III: Distribution of DoQ-verified resolvers per ASN as of July 10-16, 2023. nextDNS has the highest number of DoQ-verified resolvers deployed worldwide.

ASN	*DoQ-verified*
nextDNS (AS34939)	522 (32.91%)
ORACLE-BMC-31898 (AS31898)	149 (9.39%)
TENCENT-NET-AP (AS45090)	63 (3.97%)
DIGITALOCEAN-ASN (AS14061)	49 (3.09%)
OVH (AS16276)	47 (2.96%)
AMAZON-02 (AS16509)	35 (2.21%)
AKAMAI-LINODE-AP (AS63949)	29 (1.83%)
AS-CHOOPA (AS20473)	28 (1.77%)
ALIBABA-CN-NET (AS37963)	27 (1.70%)
CHINANET-BACKBONE (AS4134)	21 (1.32%)
AS-COLOCROSSING (AS36352)	17 (1.07%)
MNGTNET (AS199274)	16 (1.01%)
TENCENT-NET-AP-CN (AS132203)	15 (0.95%)
MICROSOFT-CORP-MSN-AS-BLOCK (AS8075)	15 (0.95%)
AS-ANEXIA (AS42473)	14 (0.88%)
ALIBABA-CN-NET (AS45102)	12 (0.76%)
MULTA-ASN1 (AS35916)	12 (0.76%)
netcup-AS (AS197540)	11 (0.69%)
IT7NET (AS25820)	11 (0.69%)
HINET (AS3462)	11 (0.69%)

continuously verified from 2022-W24 until 2023-W20 (light blue bars). Additionally, we find that only 125 (7.75%) of the initial 844 resolvers (2021–W28) are still verified in 2023-W28. However, we also notice that a steadily high number (≈1500) of new resolvers (compared to the first week of study) support DoQ (green bar, Fig. 6b) from 2022-W32 until the end of our study. This implies the rising popularity of DoQ and its implementation within the operations community. This fluctuation of DoQ reflects the development process: While DoQ was being standardized, implementations and services changed frequently over time and then stabilised after its standardization in May, 2022.

Distribution of DoQ-verified resolvers in censored or surveillance countries: We further analyzed the distribution of the *DoQ-verified* resolvers for the last week of our study (i.e. 2023-28), especially in censored or surveillance countries (see: Table II). We include this study to assess whether encrypted DNS protocols like DoQ appear to be deployed or reachable in countries with more restrictive Internet environments. We found that *DoQ-verified* resolvers are present in twelve such censored or surveillance countries. Out of them China has the most number of resolvers, whereas the United Arab Emirates, Kazakhstan and Belarus have the least (only 1). When looking at the network types of these resolvers, we find that most of the resolvers are within content delivery networks (>57%), followed by ISPs (≈22%). The least number of resolvers belong to other network types, such as enterprise, education or non-profit organizations. We emphasize that our measurements do not assess blocking or interference; they only reflect whether DoQ-verified resolvers responded to our queries.

Even before DoQs standardization, both *AdGuard* [25] and *nextDNS* [26] actually did use DoQ in production systems for their DNS-based ad and tracker blocking services, offering publicly reachable DoQ servers as well as client implementations [27], [28]. This is reflected in the *Common Names* of the X.509 certificates offered by the verified DoQ resolvers (see: Fig. 7): In 2021–W28, 182 resolvers (21.56%, golden yellow bar) state `dns.nextdns.io` as their common name.

Analyzing the change over time, we observe that *nextDNS* operates the highest share of resolvers in each week, with a mean of roughly 180 resolvers in 2021-W27 to 2021-W31, increasing to a mean of 198 resolvers in 2021-W32 to 2022-W12. The mean then slightly decreases to 189 in 2022–W16 to 2022–W24, followed by a steep decrease to a mean of roughly 56 in 2022–W28 to 2023–W16, and finally to a mean of 30 in 2023–W20 to 2023–W32. While the increase was observed between 2021-W31 and 2021-W32, *nextDNS* offered `DoQ Draft 02/QUIC Draft 32` (see: Fig. 5, yellow bars) until 2021-W32 and downgraded all resolvers to `DoQ Draft 02/QUIC Draft 29` (see: Fig. 5, green bars) in 2021-W33, where this DoQ/QUIC pair is exclusively offered by *nextDNS*. After adding support for QUIC version 1 in 2021-W43, we also observe that all *nextDNS* resolvers offer `DoQ Draft 02/QUIC 1` (see: Fig. 5, red bars) since that week; hence, we attribute the previously observed downgrade to the missing support of QUIC version 1 in our tooling during that timeframe. Finally, support for *DoQ RFC/QUIC 1* (see: Fig. 5, light green bars) was added in 2022–W13 and *nextDNS* started its support from the same time as well. Considering the publicly reachable DoQ servers of *AdGuard* (identified by the common names `dns.adguard.com` and `adguard.ch`), we identify 25 resolvers (see: Fig. 7, blue bars) offering `DoQ Draft 03/QUIC 1` (see: Fig. 5, brown bars) in 2022-W04 (2.1%) and 17 resolvers (see: Fig. 7, blue bars) offering support in 2023–W28 (1.05%). Note that this DoQ/QUIC pair is exclusively offered by the *AdGuard* services, as it differs from the *AdGuard Home (AGH)* open source DNS server implementation detailed above. We find 12–17 resolvers (see: Fig. 7, blue bars) with the common name `dns.adguard.com` and `DoQ Draft 02/QUIC Draft 34` (see: Fig. 5, blue bars) until 2021-W47, after which these resolvers switch to `DoQ Draft 03/QUIC 1` (see: Fig. 5, brown bars) starting 2021-W48. Moreover, `DoQ Draft 03/QUIC 1` is also offered by 6–8 resolvers (see: Fig. 7, green bars) using `adguard.ch` starting 2021-W49. Also, *AdGuard* added support for *DoQ RFC/QUIC 1* (see: Fig. 5, light green bars) in 2022–W21, offering 16 resolvers (see: Fig. 7, blue bars) with the common name `dns.adguard.com` and 8 resolvers (see: Fig. 7, green bars) using `adguard.ch`.

This article has been accepted for publication in IEEE Transactions on Network and Service Management. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TNSM.2026.3688441

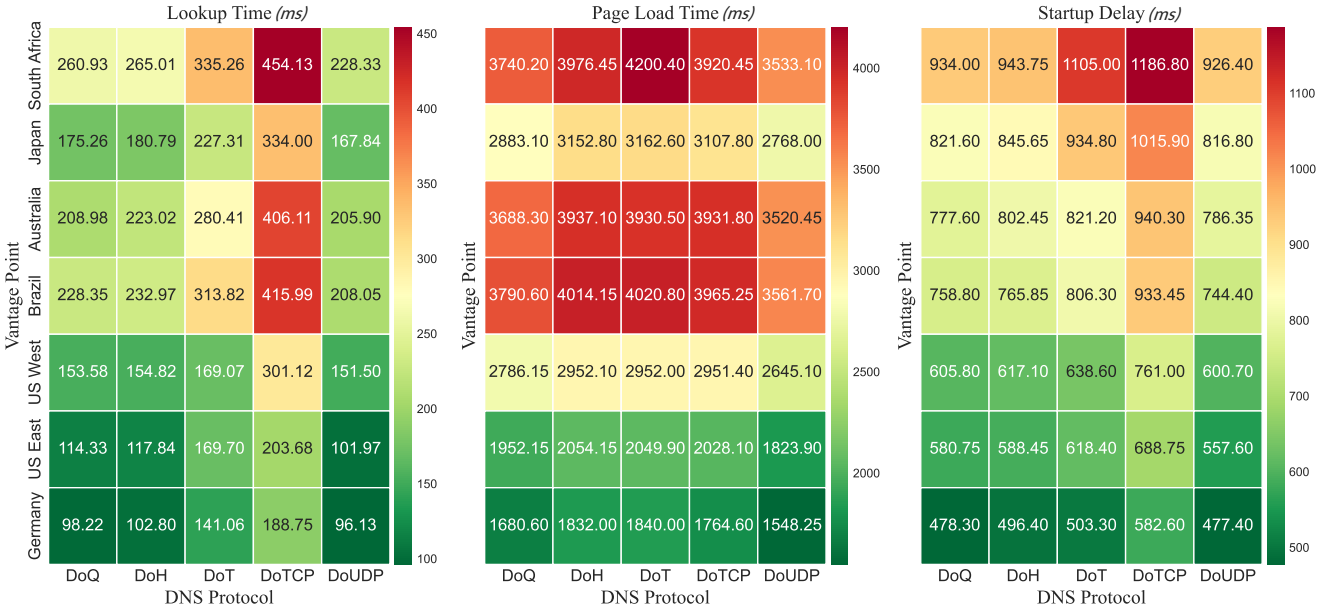

Fig. 8: Heat map showing median values (in ms) for name lookup time, PLT and startup delays for Video ID *aqz-KE-bpKQ* [52] across all resolvers. Grouped by vantage points and DNS protocols. For all metrics and protocols the vantage point in South Africa performs the worst while the ones in Europe and North America perform the best. The US West Coast appears to perform slightly worse than its East Coast. Amongst the encrypted DNS protocols, DoQ performs the best, followed by DoH.

Finally, note that `dns.nextdns.io` is present in 2023-W24 with a negligible count; the missing/near-zero value is due to aggregation and scale effects rather than the absence of resolvers.

Adoption of Encrypted DNS in Real-World: We show the distribution of DoQ-verified resolvers per ASN for the last week of our study (i.e. 2023-W28) in Table III. Similar to the observations made in Fig. 7, Table III also shows nextDNS having the highest share (32.9%) of DoQ-verified resolvers deployed worldwide.

> ***Takeaway:*** *Over the course of 2 years, we observe a steady increase with eventual saturation in adoption of DoQ by public DNS resolvers. Even when analysing the churn, we notice a steady number ($>$ 1000) of resolvers continue to support DoQ, demonstrating its growing popularity within the operations community.*

V. Impact of DoQ on YouTube Streaming

In this §, we evaluate the effect of DoQ for YouTube video streaming in comparison to DoUDP, DoTCP, DoT, and DoH. Specifically, we evaluate the performance using the metrics defined in § III-C. The evaluated metrics capture different stages of the workflow—DNS resolution, page loading, and media playback—and are sequential but not directly coupled.

Although DNS lookups contribute only a small fraction of total streaming time, prior studies [63], [64], [65] have shown that DNS has an indirect but potentially larger impact on video QoE by influencing CDN mapping and replica selection. Since DNS responses can redirect clients to different replicas (or front-end servers), suboptimal DNS resolution may steer the client to a suboptimal replica with higher latency or lower throughput. This can worsen PLT and startup delay beyond the direct DNS lookup time. At the same time, YouTube typically uses persistent HTTP/2 or HTTP/3 connections during playback, so DNS queries do not occur repeatedly per segment. Thus, the impact of DNS is concentrated in session initialization and replica selection, motivating our choice of lookup time, PLT, and startup delay as evaluation metrics. We focus on empirical evaluation under real-world conditions, where factors such as CDN selection and network variability are difficult to model analytically. Adaptive bitrate (ABR) mechanisms operate after playback begins over the established connection and are therefore outside the scope of our DNS transport analysis.

A. Performance across Different Vantage Point and Protocol Combinations

We conducted a study across different vantage points and protocol combinations, which naturally differ in latency, network paths, congestion levels, and access performance, providing heterogeneity for evaluating DNS protocol performance. This study determines the relative percentage of measurements having initial video quality of *hd720* and *large* (i.e. 854x480). We observed that for most vantage points, the distribution of video quality in our dataset is more inclined towards 720p (for $>$97% of results) hence, we restrict the rest of our evaluation to 720p.

Fig. 8 shows the median values covering these metrics, for each *(vantage point, protocol)* combination across all resolvers for video with ID *aqz-KE-bpKQ* [52]. The number of successful measurements performed for each vantage point is in the order: Brazil (16932), US West (16849), Germany (16792), Australia (16759), Japan (16748), US East (16623) and South

TABLE IV: Contribution (%) of DNS resolution time in PLT and startup delay across the median for each *[vantage point,protocol]* tuple. The impact of lookup time on PLT is lower than its impact on startup delay.

Vantage Points	Lookup Time/PLT (%)					Lookup Time/Startup Delay (%)				
	DoQ	DoH	DoT	DoTCP	DoUDP	DoQ	DoH	DoT	DoTCP	DoUDP
South Africa	6.98	6.67	7.98	11.58	6.46	27.94	28.08	30.34	38.27	24.65
Japan	6.08	5.73	7.19	10.75	6.06	21.33	22.01	24.32	32.88	20.55
Australia	5.67	5.67	7.13	10.33	5.85	26.88	27.79	34.15	43.19	26.18
Brazil	6.02	5.8	7.8	10.49	5.84	30.09	30.42	38.92	44.57	27.95
US West	5.51	5.24	5.73	10.2	5.73	25.35	25.09	26.48	39.57	25.22
US East	5.86	5.74	8.28	10.04	5.59	19.69	20.03	27.44	29.57	18.29
Germany	5.84	5.61	7.67	10.7	6.21	20.54	20.71	28.03	32.40	20.14

Africa (15918). Each metric has a separate color scale and is represented in milliseconds. The order of the vantage point is determined by median startup delay across protocols. This order is slightly different for lookup times and PLTs. It is observed that South Africa performs worst, followed by Brazil, Australia and Japan, with the best performing three staying the same across metrics.

Name lookup times: Within a vantage point, DNS protocols perform differently depending on the metric. For name lookup times, DoTCP performs the worst across vantage points. This is due to the handshake round-trip that is added for every single query as DoTCP does not implement connection reuse. The fastest lookup times are achieved with DoUDP, which is expected due to no connection setup overhead. However, as DoUDP is unencrypted, it does not feature in-network privacy. The second best protocol is DoQ, followed by DoH and lastly DoT. The differences between them highly depend on the vantage points. For example, on the German and US West vantage point, DoQ performs very well and is only around 2 ms worse than DoUDP while DoH is around 4.6 ms worse than DoQ. In contrast, for the vantage point in South Africa, DoQ takes 32.6 ms longer than DoUDP, while DoH takes further 4.1 ms longer than DoQ. The lookup time differences between vantage points are likely due to the network topology (i.e. distance from the resolver servers). Overall, for Germany, US West and Australia, the lookup time for DoQ is ≈1.5% higher than DoUDP. When comparing them using a paired t-test, we observe a pvalue = 0.0183 ($<$0.05), which implies that the results are statistically significant. We also compute the contribution of lookup time towards PLT and startup delay (Table IV). We observe that there is tangible impact of lookup time on both PLT and startup delay, while impact on PLT ($5.5\% - 11.5\%$) is lower compared to the impact on startup delay ($18.3\% - 44.6\%$). We next discuss PLT and startup delay.

Page Load Times: The PLTs exhibit performance similar to that of lookups, but here differences between the vantage points are accentuated. As aforementioned: vantage points in South Africa, Brazil and Australia perform the worst. There are likely multiple factors influencing this observation but the main cause is the distance between the YouTube servers and the resolvers. DoT and DoTCP also perform worse compared to other protocols due to aforementioned reasons. PLT evaluations includes at least six DNS resolutions. Some of them likely happen in parallel, triggering the DoT pool implementation to cause more handshakes. If a TLS connection is established for a DNS query, and subsequently another query is sent to the DNS proxy, it creates a new connection and runs a new handshake. As a result, multiple TLS connections are created for one DoT instance of the DNS proxy, which is then referred to as the DoT pool. To confirm, the number of TLS handshakes that are triggered before the page onload event, every DoT measurement run is calculated. It is observed that 76.8%, 9.8%, 8.8% and 2.7% of the measurements have three, four, two and one handshake per measurement run respectively. Explicitly, DoQ has exactly one handshake for all measurements, while DoTCP has six, eight and seven handshakes for 91.1%, 3.1% and 2.2% of the measurements respectively. As such, these additional handshakes cause DoT and DoTCP to perform worse compared to other DoX protocols across all metrics.

In summary, the observed differences across vantage points are likely influenced by factors such as geographic proximity to CDN replicas, variations in RTT, and underlying routing characteristics.

> ***Takeaway:*** *For Germany, US West and Australia, the lookup time for DoQ is ≈1.5% higher than DoUDP, whereas for DoH it is ≈7% higher than DoUDP. On the contrary, for PLT, the difference in results for the same vantage points are more pronounced where DoQ performs ≈5.5% worse compared to DoUDP. However, PLT as a metric has less impact on video streaming, so this is less of an issue for video workloads.*

Startup Delay: In line with the previous observations, DoTCP has the highest startup delay. The magnitude of this delay is indicative of how DNS can play a substantial role in the initial startup of the video. The throughput achieved between the vantage points and the YouTube servers is also a contributing factor here. This implies that the difference between the startup delays of the various DoX protocols likely does not scale proportionally with the handshake round-trip differences between them. On the vantage point in Germany, for instance, DoTCP performs $105.2\ ms$ worse than DoUDP which increases to $260.4\ ms$ at the vantage point in South Africa. An additional

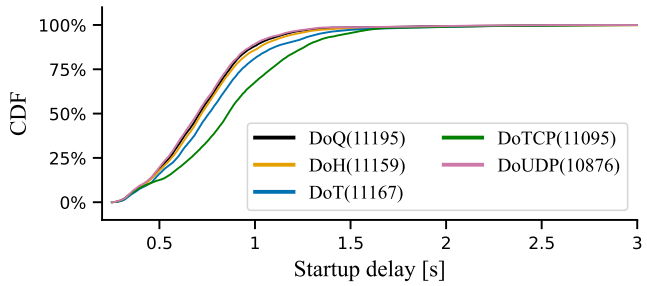

(a) CDF of startup delays over all vantage points. The markers indicate percentiles in steps of 10%.

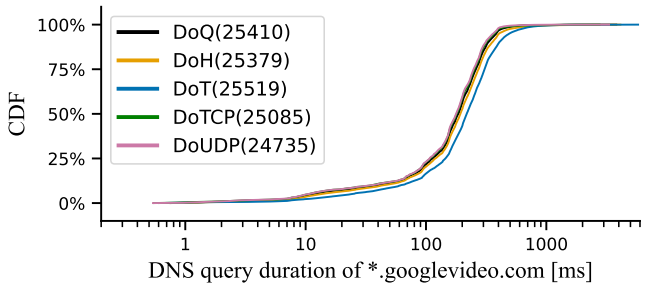

(b) DNS Query durations (i.e., the DNS lookup duration without any handshakes) over all vantage points for all resolutions of `googlevideo`.

Fig. 9: CDFs of the startup delays (a, top) and Query durations (b, bottom) over all vantage points. Among the encrypted DNS protocols, DoQ shows the fastest startup times while DoT is considerably slower (see a). We attribute the slow DoT startup times to the inflated DoT query durations (see b).

observation to note is that the median startup delay for DoT appears to be worse than its corresponding DoH values.

To dig deeper, we investigate the timeline of when DNS resolutions occur within a measurement in further detail. The video playback starts strictly after the page loads, while the video ID is set already once the video element has been located on the website. This suggests that the PLT is guaranteed to contain the DNS protocol's (amortized) handshake duration whereas the startup delay starts strictly after the initial DNS handshake has finished. The exceptions to this observation are DoTCP (no connection reuse) and DoT (see below). However, domain names of video chunks (`*.googlevideo.com`) are likely resolved only at the very outset. If different domains are used during playback they are unlikely resolved in parallel and initial page load would have already created multiple pooled connections (as explained above). We observed around 353 measurements (out of 23,517 DoT measurements) where at least one DoT handshake is recorded after the PLT whereas only 42 measurements exist in the dataset where a DoT handshake is initiated after the first video chunk domain is resolved. The number of measurements where this occurs is relatively small compared to the full sample size. Consequently, the data shows a skewed median startup delay for DoT across all vantage points.

Fig. 9a depicts the cumulative distribution of startup delay across all vantage points. Given that DoTCP does not offer encryption and is also meant as a fallback mechanism for DoUDP, we do not consider DoTCP in further analysis. It is observed that DoUDP has a median startup delay of 698.9 ms which increases to 1012.8 ms for the 90^{th} percentile. The second best on average is DoQ with 708.6 ms (1031.5 ms at 90^{th} percentile), followed by DoH with 720.8 ms (1066.0 ms at 90^{th} percentile). The worst performing one is DoT with 758.6 ms (1176.7 ms at 90^{th} percentile). The plot confirms that DoT appears to be worse at all percentiles, i.e., across all vantage points and for most measurements. To perform a causal analysis, we investigate the captured DNS metrics in Fig. 9b. We find that query duration (i.e., the lookup duration without any handshakes) appears to be the cause of the anomaly with DoT. Since the query duration is inflated, lookup times are increased, which then affects the startup delay. This is visible across all vantage points, for all resolvers regardless of initial quality. The distribution for query durations is shown in Fig. 9b. We observe that median value for DoT is 217.5 ms, while for DoH it is 193 ms. DoQ has a median query duration of 185.7 ms and DoUDP of 177.5 ms. Comparing this to the query duration for `www.youtube.com`, the medians are: DoQ 175.1 ms, DoUDP 173 ms and DoT 173.8 ms. Thus, performance of the DoE protocols follow the order: DoQ $<$ DoH $<$ DoT (where less is better.)

Table V shows the best and worst case startup delay across *(vantage point, protocol)* combinations while the corresponding median values were already depicted in Fig 8. We evaluate individual vantage points in more detail. For instance, with the vantage point in South Africa, we observe that differences between DoUDP, DoQ and DoH are small and the distributions are quite similar. On the contrary, in case of DoT, the startup delay increases comparatively more due to aforementioned reasons. Correspondingly, for the vantage point in Germany, we observe that differences between protocols (including DoTCP) are less pronounced up to roughly the 25^{th} percentile. This is probably due to low round-trip times to recursive resolvers, YouTube servers, or a combination of both. The resolvers do have a location bias toward Asia, Europe and North America, as is also seen in previous studies [23]. Even though, the expected performance differences are visible at the median, the 90^{th} percentile shows the differences more prominently. The next vantage point examined in detail is US West. Here, DNS protocols perform similarly up to the 10^{th} percentile. The median startup delay is slightly higher compared to German vantage point, but the 90^{th} percentile is lower, indicating that startup delays are more stable for all DNS protocols. This vantage point is also the one with the lowest 90^{th} percentile value compared to other vantage points. This vantage point was found to have the most number of measurements where at least one DoT handshake occurred after the first `googlevideo` domain name resolution. However, difference between DoT and DoH is larger on other vantage points which means that additional handshakes are not the root cause. The last vantage point shown in detail is the one in Japan where the differences observed between DoT and the other protocols are very high, albeit not as high as in South Africa. Observing only DoH and DoT, which in theory should have almost exactly the same startup delays, it is observed that DoT has a median increase of around 89.1 ms while the 90^{th} percentile increase is 79 ms over DoH.

In summary, we observe that DoTCP performs worst across

TABLE V: P10 and P90 values startup delays accross vantage point and protocol combinations. DoTCP performs worst whereas DoQ performs best amongst the encrypted DNS protocols generally across all vantage points for both best and worst-case scenarios. The worst performing results are shown in yellow whereas the best ones amongst the DoX protocols are shown in violet. The only exception where DoQ performs better than DoUDP is highlighted in pink.

DNS Protocols	Startup Delay (ms)													
	Germany		US East		US West		Brazil		Australia		Japan		South Africa	
	P10	P90	P10	P90	P10	P90	P10	P90	P10	P90	P10	P90	P10	P90
DoQ	313.2	848.1	430.2	863.9	605.8	761.9	564.7	987.9	541.2	968.2	479.3	1193.3	711.5	1618.6
DoH	313.3	873.7	432.9	893.8	617.1	776.4	765.9	1014.3	555.7	982.3	469.8	1238.1	716.3	1781.3
DoT	318.2	921	451.5	899.2	638.6	819.8	595.4	1076.8	571.4	1012.6	555.8	1317.1	759	1931.3
DoTCP	311.1	1066	454.4	1039.1	761	896.4	707.6	1320.1	655.7	1215.1	557.3	1461.2	859.3	1935.7
DoUDP	306.6	840.3	424.5	837.9	600.7	739.4	560.6	969.2	529.7	958.7	470.8	1152.1	715.6	1616.3

all vantage points. This is mainly because current DoTCP implementations lack connection reuse across queries. As a result, each lookup incurs a full TCP three-way handshake and is additionally affected by slow-start and head-of-line blocking, leading to consistently higher lookup times. A recent study [50] also reports significant real-world DoTCP query failures, further contributing to latency overhead, and highlights that TCP-level optimizations to reduce latency are often not adopted by DNS servers. Additionally, DoUDP performs best across all vantage points followed by DoQ. The only exception is observed for the 10^{th} percentile at the vantage point in South Africa and for the median at the vantage point in Australia where DoQ performs better than DoUDP.

Takeaway: *Across the best performing vantage points, median startup delay increase of DoQ over DoUDP is* $< 1\%$ *whereas that of DoH over DoUDP varies between* $1.8\% - 3.9\%$. *Even for the worst case scenario (i.e.* 90^{th} *percentile), DoQ shows a similar trend over DoUDP. However, the* 90^{th} *percentile results for DoH over DoUDP are relatively worse with variations between* $3.9\% - 10.2\%$. *This means that DoH degrades more strongly in the high-delay (tail) regime, whereas DoQ remains more stable.*

B. DoQ vs. other DNS Protocols

It is evident from the previous section that DoTCP does not perform well while DoT unexpectedly performs worse compared to DoH even for higher query duration under specific measurement runs. Hence, we do not discuss these protocols further. We rather focus on providing an overall view highlighting the differences between the different protocols without the influence of resolvers and their geographical location. Therefore, we now evaluate the increase (or decrease) of PLTs and startup delays to emphasize when using DoQ over DoUDP or DoH is beneficial.

For providing a macro perspective, the median value of the metric for each DNS protocol is calculated by grouping results with the same *vantage point, resolver* for the startup delay and the PLT. This allows comparing the DNS protocols to each other within one such grouping. We consider DoQ as the baseline and the relative increase or decrease over the median value with DoQ is calculated against the median values with

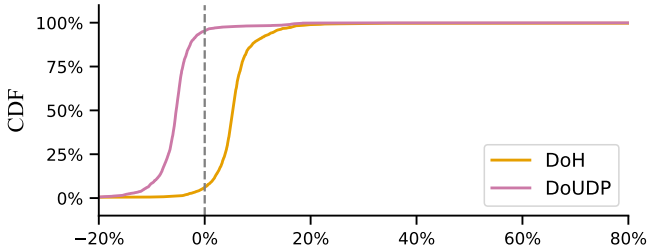

(a) Relative PLT increase.

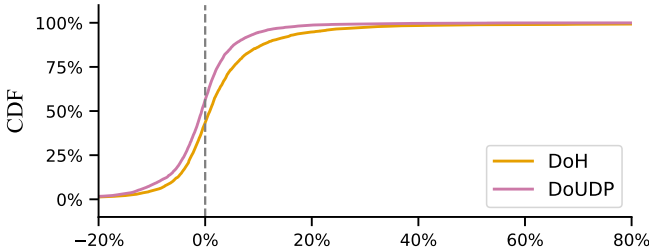

(b) Relative startup delay increase.

Fig. 10: CDFs of the relative performance increases of DoH and DoUDP compared to the DoQ baseline. Each data point represents a vantage point, resolver and DNS protocol tuple. For PLTs (a, top), DoUDP and DoH perform as expected in relation with DoQ. For startup delay (b, bottom), DoQ shows quite an unexpected result with it achieving comparable performance benefits to DoUDP from the median and eventually excelling over DoUDP from the 60^{th} percentile.

DoUDP and DoH. This can be seen as a form of normalizing the data to a baseline protocol, i.e., DoQ.

Page Load Time: The results shown in Fig. 10a depict differences between the protocols where DoUDP appears to perform the best overall. The median relative decrease in PLT when using DoUDP compared to the DoQ baseline is -5.4% while the 10^{th} and the 90^{th} percentile records a decrease in PLT of -9.4% and -2.3% respectively. Thus, for a majority of the results DoUDP is strictly better when it comes to PLTs. DoH on the other hand is worse as P10, P50 and P90 records a relative increase of 1.5%, 5.3% and 10.0% respectively.

Startup Delay: The normalized startup delay is shown in Fig. 10b. It is observed that the relative changes between the three protocols are less pronounced. This is apparent as the DNS protocols do not set up new connections at this point of the measurement run and thus lookup duration is

minimal, which then results in few differences between the startup delays. By comparing DoUDP to the DoQ baseline, the P10 relative change is observed to be -8.8% which increases to -0.7% for the median. This means at the median, there is no discernible difference between DoQ and DoUDP when it comes to video startup. However, for P90 this increases to 6.6%. A similar observation is recorded for DoH where P10 is 6.2%, the median is 0.9% and P90 is 12.6%. Overall, for PLTs there is still a small difference between the protocols that follows the order (DoUDP $>$ DoQ $>$ DoH), however on average there is little difference.

Takeaway: *When vantage points and resolvers are normalized, DoQ has quite a similar performance to DoUDP for the median startup delay. The results for DoQ eventually starts improving with almost* 40% *of our samples performing better than DoUDP. As startup delay is a more relevant metric in case of YouTube video streaming, therefore, enabling DoQ at the cost of minimal overheads in terms of PLT compared to DoUDP is the best choice to preserve in-network privacy.*

VI. Conclusion and Future Work

To promote real-world adoption of DNS over QUIC (DoQ), we conducted the first comprehensive, long-term evaluation of its deployment and performance during YouTube video streaming. Our 25-month longitudinal analysis shows a steady rise in DoQ adoption by public DNS resolvers, eventually reaching a plateau, reflecting its gradual integration into the DNS ecosystem. We developed a robust measurement methodology to compare five DNS protocols using seven globally distributed vantage points and a curated set of 312 DoX-verified resolvers. Our findings indicate that DoQ consistently outperforms other encrypted DNS protocols in lookup time and PLT. In particular, the median increase in the startup delay of DoQ is less than $< 1\%$ over DoUDP, with this marginal difference persisting up to the 90^{th} percentile. Remarkably, DoQ achieves lower startup delays than DoUDP in approximately 40% of the measured cases. These results highlight that DoQ not only offers the best performance among encrypted DNS protocols but also incurs minimal overhead compared to the baseline unencrypted protocol, DoUDP. In conclusion, our study positions DoQ as a highly effective and efficient choice for DNS resolution in latency-sensitive applications such as video streaming.

A limitation of our study is that we restrict video experiments to YouTube and two ad-free videos only to ensure global repeatability. Future work will extend the analysis to platforms with different architectures to evaluate whether our findings generalize across video ecosystems. Another limitation of our setup is that client hardware, browser configuration, and access-link characteristics remain fixed across all vantage points. Future work will examine whether factors such as device types, mobile vs. wired networks, and bandwidth variability amplify or mitigate DoQ's performance benefits. Lastly, incorporating path analysis alongside transport measurements is a valuable direction for future work and would enable deeper investigation into the impact of routing and geographic distance on encrypted DNS performance.

References

[1] C. Deccio and J. Davis, "DNS Privacy in Practice and Preparation," in *Conference on Emerging Networking Experiments And Technologies (CoNEXT)*, 2019, pp. 138–143.

[2] D. W. Kim and J. Zhang, "You Are How You Query: Deriving Behavioral Fingerprints from DNS Traffic," in *Security and Privacy in Communication Networks (SecureComm)*, vol. 164, 2015, pp. 348–366.

[3] M. Kirchler, D. Herrmann, J. Lindemann, and M. Kloft, "Tracked Without a Trace: Linking Sessions of Users by Unsupervised Learning of Patterns in Their DNS Traffic," in *Proceedings of the 2016 ACM Workshop on Artificial Intelligence and Security, CCS*, 2016, pp. 23–34.

[4] J. Li, X. Ma, G. Li, X. Luo, J. Zhang, W. Li, and X. Guan, "Can We Learn what People are Doing from Raw DNS Queries?" in *IEEE International Conference on Computer Communications (INFOCOM)*, 2018, pp. 2240–2248.

[5] S. Da Silva, S. Ben Mokhtar, S. Contiu, D. Négru, L. Réveillère, and E. Rivière, "PrivaTube: Privacy-Preserving Edge-Assisted Video Streaming," in *Proceedings of the 20th International Middleware Conference*, 2019, p. 189–201.

[6] 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 (PAM)*, 2021, pp. 192–209.

[7] A. Hounsel *et al.*, "Can Encrypted DNS Be Fast?" in *PAM*, 2021.

[8] T. Böttger, F. Cuadrado, G. Antichi, E. L. a. Fernandes, G. Tyson, I. Castro, and S. Uhlig, "An Empirical Study of the Cost of DNS-over-HTTPS," in *Proceedings of the Internet Measurement Conference*, 2019, pp. 15–21.

[9] 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 *Internet Measurement Conference*, 2019, pp. 22–35.

[10] R. Chhabra *et al.*, "Measuring DNS-over-HTTPS Performance around the World," in *IMC*, 2021.

[11] A. Hounsel, K. Borgolte, P. Schmitt, J. Holland, and N. Feamster, "Comparing the Effects of DNS, DoT, and DoH on Web Performance," in *Proceedings of The Web Conference*, 2020, pp. 562–572.

[12] K. Borgolte, T. Chattopadhyay, N. Feamster, M. Kshirsagar, J. Holland, A. Hounsel, and P. Schmitt, "How DNS over HTTPS is reshaping Privacy, Performance, and Policy in the Internet Ecosystem," *Performance, and Policy in the Internet Ecosystem*, 2019.

[13] J. Iyengar and M. Thomson, "QUIC: A UDP-Based Multiplexed and Secure Transport," RFC 9000, May 2021.

[14] M. Thomson and S. Turner, "Using TLS to Secure QUIC," RFC 9001, May 2021.

[15] J. Iyengar and I. Swett, "QUIC Loss Detection and Congestion Control," RFC 9002, May 2021.

[16] C. Huitema, S. Dickinson, and A. Mankin, "DNS over Dedicated QUIC Connections," RFC 9250, May 2022.

[17] K. Buchholz, "YouTube Responsible For 16% Of Global Internet Traffic," https://www.forbes.com/sites/katharinabuchholz/2025/04/25/youtube-responsible-for-16-of-global-internet-traffic/, 2025, Accessed: 09.12.2025.

[18] D. Ruby, "YouTube Statistics 2023: Data For Brands & Creators," https://www.demandsage.com/youtube-stats/, 2023, Accessed: 05.05.2023.

[19] S. Aslam, "YouTube by the Numbers: Stats, Demographics & Fun Facts," https://www.omnicoreagency.com/youtube-statistics/, 2023, Accessed: 02.05.2023.

[20] D. K. Krishnappa, D. Bhat, and M. Zink, "DASHing YouTube: An analysis of using DASH in YouTube video service," in *38th Annual IEEE Conference on Local Computer Networks*, 2013, pp. 407–415.

[21] Z. Hu, L. Zhu, J. Heidemann, A. Mankin, D. Wessels, and P. E. Hoffman, "Specification for DNS over Transport Layer Security (TLS)," RFC 7858, May 2016.

[22] P. E. Hoffman and P. McManus, "DNS Queries over HTTPS (DoH)," RFC 8484, Oct. 2018.

[23] M. Kosek *et al.*, "One to Rule Them All? A First Look at DNS over QUIC," in *PAM 2022*, 2022.

[24] M. Kosek, L. Schumann, R. Marx, T. V. Doan, and V. Bajpai, "DNS Privacy with Speed? Evaluating DNS over QUIC and Its Impact on Web Performance," in *Proceedings of the 22nd ACM Internet Measurement Conference*, 2022, p. 44–50.

[25] "AdGuard DNS-over-QUIC," Accessed 2023-Feb-21. [Online]. Available: https://adguard.com/en/blog/dns-over-quic.html

[26] "NextDNS Knowledge Base," Accessed: 31.07.2023]. [Online]. Available: https://help.nextdns.io/t/x2hmvas/what-is-dns-over-tls-dot-dns-over-quic-doq-and-dns-over-https-doh-doh3

[27] "DNS Lookup," Accessed 2023-Jun-30. [Online]. Available: https://github.com/ameshkov/dnslookup

[28] "NextDNS CLI Client," Accessed 2023-Jun-30. [Online]. Available: https://github.com/nextdns/nextdns

[29] 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.

[30] J. Sengupta, M. Kosek, J. Fries, P. Dikshit, and V. Bajpai, "Web Privacy By Design: Evaluating Cross-layer Interactions of QUIC, DNS and H/3," in *2023 IFIP Networking Conference (IFIP Networking)*, 2023, pp. 1–9.

[31] J. Sengupta, M. Kosek, J. Fries, S. Ferlin, P. Dikshit, and V. Bajpai, "On Cross-Layer Interactions of QUIC, Encrypted DNS and HTTP/3: Design, Evaluation and Dataset," *IEEE Transactions on Network and Service Management*, vol. 21, no. 3, pp. 2992–3007, 2024.

[32] M. Lyu, H. H. Gharakheili, and V. Sivaraman, "A Survey on DNS Encryption: Current Development, Malware Misuse, and Inference Techniques," *ACM Comput. Surv.*, vol. 55, no. 8, dec 2022.

[33] G. Kambourakis and G. Karopoulos, "Encrypted DNS: The Good, the Bad and the Moot," *Computer Fraud & Security*, vol. 2022, no. 5, 2022.

[34] P. Bielefeld, F. Hoffmann, S. Sassalla, V. Ververis, and V. Bajpai, "The Future of DNS Privacy: A Comparison of DNS over QUIC and DNS over HTTP/3," in *Passive and Active Measurement: 27th International Conference, PAM 2026*, 2026, p. 202–228.

[35] S. Siby, M. Juárez, C. Díaz, N. Vallina-Rodriguez, and C. Troncoso, "Encrypted DNS -> Privacy? A Traffic Analysis Perspective," in *27th Annual Network and Distributed System Security Symposium*, 2020.

[36] M. Wang, A. Kulshrestha, L. Wang, and P. Mittal, "Leveraging strategic connection migration-powered traffic splitting for privacy," *Proc. Priv. Enhancing Technol.*, vol. 2022, no. 3, pp. 498–515, 2022.

[37] G. Hu and K. Fukuda, "An Analysis of Privacy Leakage in DoQ Traffic," in *Proceedings of the CoNEXT Student Workshop*, 2021, p. 7–8.

[38] M. A. Rajan, A. Varghese, N. Narendra, M. S. D. Singh, V. L. Shivraj, M. G. Chandra, and P. Balamuralidhar, "Security and Privacy for Real Time Video Streaming Using Hierarchical Inner Product Encryption Based Publish-Subscribe Architecture," in *Conference on Advanced Information Networking and Applications Workshops*, 2016.

[39] H. M. Moghaddam, G. Acar, B. Burgess, A. Mathur, D. Y. Huang, N. Feamster, E. W. Felten, P. Mittal, and A. Narayanan, "Watching You Watch: The Tracking Ecosystem of Over-the-Top TV Streaming Devices," in *Conference on Computer and Communications Security, (CCS)*, 2019, pp. 131–147.

[40] F. Li, J. W. Chung, and M. Claypool, "Silhouette: Identifying YouTube Video Flows from Encrypted Traffic," in *Workshop on Network and Operating Systems Support for Digital Audio and Video, NOSSDAV*, 2018, pp. 19–24.

[41] Y. Wang, M. Lyu, and V. Sivaraman, "Characterizing User Platforms for Video Streaming in Broadband Networks," in *Proceedings of the 2024 ACM on Internet Measurement Conference*, 2024, p. 563–579.

[42] M. Lyu, H. Habibi Gharakheili, C. Russell, and V. Sivaraman, "Enterprise DNS Asset Mapping and Cyber-Health Tracking via Passive Traffic Analysis," *IEEE Transactions on Network and Service Management*, vol. 20, no. 3, pp. 3699–3716, 2023.

[43] "The ZMap Project," Accessed 2023-Jan-31. [Online]. Available: https://zmap.io/

[44] ZMap, "UDP Data Probes," Accessed 2023-Jan-31. [Online]. Available: https://github.com/zmap/zmap/blob/master/examples/udp-probes/README

[45] "ZMap DoQ," Accessed 2023-Feb-21. [Online]. Available: https://github.com/mgranderath/zmap-doq

[46] J. Rüth, I. Poese, C. Dietzel, and O. Hohlfeld, "A First Look at QUIC in the Wild," in *Passive and Active Measurement*, 2018, pp. 255–268.

[47] "Verify DoQ," Accessed 2023-Jan-31. [Online]. Available: https://github.com/mgranderath/verify-doq

[48] "Misc DNS Measurements," Accessed 2023-Jun-30. [Online]. Available: https://github.com/mgranderath/misc-dns-measurements

[49] "DNSPerf," Accessed 2023-Jun-30. [Online]. Available: https://github.com/mgranderath/dnsperf

[50] M. Kosek *et al.*, "Measuring DNS over TCP in the Era of Increasing DNS Response Sizes: A View from the Edge," *SIGCOMM CCR*, vol. 52, no. 2, pp. 44–55, June 2022.

[51] "IP Geolocation API," Accessed 2025-12-10. [Online]. Available: https://ip-api.com/

[52] B. Institute, "Big Buck Bunny 60fps 4K," https://www.youtube.com/watch?v=aqz-KE-bpKQ, 2014, accessed on 24.05.2023.

[53] ——, "Glass Half - Blender Animated Cartoon," https://www.youtube.com/watch?v=lqiN98z6Dak, 2015, accessed on 24.05.2023.

[54] A. Schwind, M. Seufert, O. Alay, P. Casas, P. Tran-Gia, and F. Wamser, "Concept and implementation of video QoE measurements in a mobile broadband testbed," in *2017 Network Traffic Measurement and Analysis Conference (TMA)*, 2017.

[55] A. Schwind, C. Midoglu, O. Alay, C. Griwodz, and F. Wamser, "Dissecting the performance of YouTube video streaming in mobile networks," *International Journal of Network Management*, vol. 30, no. 3, 2020.

[56] YouTube, "YouTube Player API Reference for iframe Embeds," 2022, Accessed 2023-Aug-17. [Online]. Available: https://developers.google.com/youtube/iframe_api_reference

[57] Z. Durumeric, E. Wustrow, and J. A. Halderman, "ZMap: Fast Internet-wide Scanning and Its Security Applications," in *22nd USENIX Security Symposium* , Aug. 2013, pp. 605–620.

[58] C. Partridge and M. Allman, "Ethical Considerations in Network Measurement Papers," *Commun. ACM*, vol. 59, no. 10, pp. 58–64, 2016.

[59] I. R. Learmonth, G. Grover, and M. Knodel, "Guidelines for Performing Safe Measurement on the Internet," IETF, Internet-Draft draft-irtf-pearg-safe-internet-measurement-05, Jul. 2021. [Online]. Available: https://datatracker.ietf.org/doc/draft-irtf-pearg-safe-internet-measurement/

[60] V. Bajpai, A. Brunstrom, A. Feldmann, W. Kellerer, A. Pras, H. Schulzrinne, G. Smaragdakis, M. Wählisch, and K. Wehrle, "The Dagstuhl Beginners Guide to Reproducibility for Experimental Networking Research," *SIGCOMM Comput. Commun. Rev.*, vol. 49, pp. 24–30, Feb. 2019.

[61] "AdGuard Home," Accessed 2023-Jun-30. [Online]. Available: https://github.com/AdguardTeam/AdguardHome

[62] "AdGuard Home Release 0.107.0," Accessed 2023-Jan-31. [Online]. Available: https://github.com/AdguardTeam/AdGuardHome/releases/tag/v0.107.0

[63] B. Ager, W. Mühlbauer, G. Smaragdakis, and S. Uhlig, "Comparing DNS resolvers in the wild," in *Proceedings of the 10th ACM SIGCOMM Internet Measurement Conference, IMC*, 2010, pp. 15–21.

[64] J. P. Rula and F. E. Bustamante, "Behind the Curtain: Cellular DNS and Content Replica Selection," in *Proceedings of the 2014 Internet Measurement Conference, IMC*, 2014, pp. 59–72.

[65] H. Hours, E. W. Biersack, P. Loiseau, A. Finamore, and M. Mellia, "A study of the impact of DNS resolvers on CDN performance using a causal approach," *Computer Networks*, vol. 109, pp. 200–210, 2016.