Web Censorship Measurements of HTTP/3 over QUIC

Kathrin Elmenhorst
Bertram Schütz
Nils Aschenbruck
[kelmenhorst,schuetz,aschenbruck]@uos.de
Osnabrück University - Institute of Computer Science
Osnabrück, 49076, Germany

Simone Basso
simone@openobservatory.org
Open Observatory of Network Interference (OONI)

ABSTRACT

Web traffic censorship limits the free access to information, making it a global human rights issue. The introduction of HTTP/3 (HTTP over QUIC) yields promising expectations to counteract such interference, due to its novelty, build-in encryption, and faster connection establishment. To evaluate this hypothesis and analyze the current state of HTTP/3 blocking, we extended the open-source censorship measurement-tool OONI with an HTTP/3 module. Using an input list of possibly-blocked websites, real-world measurements with HTTPS and HTTP/3 were conducted in selected Autonomous Systems in China, Iran, India, and Kazakhstan. The presented evaluation assesses the different blocking methodologies employed for TCP/TLS versus the ones employed for QUIC. The results reveal dedicated UDP blocking in Iran and major IP blocklisting affecting QUIC in China and India.

ACM Reference Format:
Kathrin Elmenhorst, Bertram Schütz, Nils Aschenbruck, and Simone Basso. 2021. Web Censorship Measurements of HTTP/3 over QUIC. In *Proceedings of ACM IMC 2021.* ACM, New York, NY, USA, 7 pages.
https://doi.org/10.1145/3487552.3487836

1 INTRODUCTION

Internet censorship interferes with accessing selected resources and information on the Internet. In several countries, governmental censors frequently limit the access to social networks and block specific websites. This is enforced via firewalls, on-path, and off-path middle boxes. In the media and in the literature, the most prominent case of such interference is certainly the "Great Firewall" of China [6]. But, Iran [1, 4], India [19], and countries of the Former Soviet Union [17, 20] have also frequently been studied.

Over the last years, the widespread usage of HTTPS shifted censorship techniques towards interfering with TLS. The problem of measuring TLS handshake blocking based on the content of the Server Name Indication field in the Client Hello has been widely studied by many [19, 22]. In early 2021, a new version of the HTTP protocol, HTTP/3, was announced. HTTP/3 uses QUIC as the underlying encrypted transport. In contrast to traditional HTTPS over TCP, QUIC uses UDP and is implemented in user-space. QUIC provides always-on, built-in encryption and reduce connection setup time [12]. Considerations in the literature assume that QUIC connections inherit a lower vulnerability to tampering and modification by middle boxes [9, 13]. Yet, to our best knowledge, this has not been empirically evaluated before.

We aim to test these expectations and provide a first survey of the current state of HTTP/3 (and QUIC) blocking. To this end, we integrated a QUIC module in the Open Observatory of Network Interference (OONI) Probe software [8]. Then, we conducted real-world side-by-side HTTPS and HTTP/3 measurements. We measured selected likely-blocked URLs in Autonomous Systems suspected of network interference in China, Iran, India, and Kazakhstan.

Our main findings are the following. In China and India, IP blocking affects HTTPS and HTTP/3 traffic alike. In Iran, we noticed the application of different blocking methods for HTTPS and HTTP/3: HTTPS traffic is mainly filtered based on the SNI, HTTP/3 is impaired by UDP endpoint blocking.

2 ETHICAL CONSIDERATIONS

Gaining knowledge through measurements is key to the understanding of censorship systems. Publishing the results openly raises public awareness and increases pressure on censoring governments and ISPs. While censors could misuse our findings, we believe that raising public awareness outweighs potential drawbacks. This stance is in line with OONI's mission of increasing transparency of Internet censorship. Due to the political nature of the subject and risks involved, the conduct of censorship measurements should be closely accompanied by ethical considerations. First and foremost, the safety of the volunteers has to be ensured. Some participants live in countries with strict internet regulations, e.g., Iran, and risk legal persecution. Thus, all participants volunteered freely and were clearly informed about the risks. Yet, we still excluded certain categories of websites to avoid raising any flags. Websites from the following categories are removed from the set of test domains: Sex Education, Pornography, Dating, Religion and LGBTQ+.

3 BACKGROUND & RELATED WORK

The following section presents the necessary background information and summarizes the related work.

3.1 HTTP/3 over QUIC

QUIC [11] is a connection-oriented, general purpose protocol, which integrates transport layer functionality with built-in encryption on top of UDP. The protocol is implemented in user-space [12], shifting the connection management to the encrypted application

IMC '21, November 2–4, 2021, Virtual Event, USA

ACM ISBN 978-1-4503-9129-0/21/11…$15.00
https://doi.org/10.1145/3487552.3487836

layer. These characteristics make QUIC a long awaited alternative to traditional TCP, especially for web browsing. Thus, the next major version of the Hypertext Transfer Protocol, HTTP/3 [3], uses QUIC as its underlying transport protocol.

3.2 Web Traffic Censorship & Error Types

This paper focuses on Internet censorship in the form of website access blocking. Website blocking methods can be split into two categories: identification and interference (c.f., [9]): While identification refers to the way the censor detects that traffic is directed towards a blocklisted website, interference is the method of blocking or impairing such traffic.

Website censorship appears in the form of network errors, where the specific type of error triggered, depends on the applied blocking method. Censors can identify traffic flows by their header fields on the IP or transport layer, e.g., destination IP or port number. IP-based identification and filtering potentially causes collateral damage because multiple services can be hosted at the same IP address. Even more information can be revealed by Deep Packet Inspection, which filters application layer packets based on keywords in their unencrypted parts, like plain HTTP payloads, or TLS extensions, e.g., Server Name Indication (SNI). To determine the censor's identification method, a common technique in censorship measurements is to identify the last successful connection establishment step. For example, it is possible to check if the failure occurs during routing or later in the TLS handshake.

Common interference methods range from DNS manipulation, over out-of-band attacks against the connection, to middle boxes dropping unwanted packets (black holing). Connection reset errors occur, when the connection was terminated due to an injected reset packet from an outside attacker [9]. A timeout failure at one of the endpoints can be a sign for black holing, but it can also be caused by various other types of network malfunctions. By regularly repeating measurements and observing the outcome over a longer period of time it is possible to recognize censorship patterns and decrease the bias caused by temporary network disturbances unrelated to censorship. In this work, we focus only on the most common error types and their relevance for censorship, denoted by the following abbreviations:

TCP-hs-to	TCP handshake timeout
TLS-hs-to	TLS handshake timeout
QUIC-hs-to	QUIC handshake timeout
conn-reset	connection reset during TLS handshake
route-err	IP routing error

We will discuss in Section 5.1 and 5.2, how these error types are associated with certain censorship methods.

3.3 Open Observatory of Network Interference (OONI)

The community-based Open Observatory of Network Interference (OONI) [8] project develops open source client software for decentralized evaluation of Internet censorship. OONI's measurement tool (*OONI Probe*) contains multiple tests to identify several internet censorship techniques, e.g., DNS manipulation, IP and TCP endpoint blocking. Connection timeouts can be identified as well as connection termination due to injected reset packets. Participating end users deploy the software on their local machines and conduct the measurements in their respective networks. The so collected data covers over 200 countries, including all 22 countries considered *not free* by the 2020 *Freedom on the Net* report [5]. Prior to our work, only HTTP/2 measurements could be conducted. To also evaluate network interference with HTTP/3, we extended the *OONI Probe* software with a dedicated HTTP/3 module.

3.4 Related Publications

Currently, two IETF drafts refer to QUIC in the context of censorship. The QUIC Human Rights Review [13] underlines the build-in encryption of QUIC to protect against deep packet inspection. The same draft also discusses the improved robustness against connection reset attacks, which is also taken up by an IETF censorship survey [9]. Such reset attacks rely on out-of-band interference, where the censor can inspect copies of packets, instead of dropping them as they arrive. According to [13] and [9], QUIC is better protected against this vulnerability, because established QUIC connections can not be easily terminated by an outsider. A censor would have to keep up inline blocking of QUIC connections, which is very resource exhausting.

Apart from IETF working groups, there currently is very little published research dealing with QUIC in the context of censorship. Kyle Hogan [10] explores the potential of running Tor over QUIC, focusing on the potential performance gain due to the multiplexing property. Zhan et al. (2021), from the Chinese Academy of Sciences, test the vulnerability of QUIC traffic in regard to website fingerprinting based on machine learning models [23]. They suggest that pattern recognition of QUIC traffic features is feasible under well-designed conditions. Alongside OONI, there are several measurement platforms dedicated to observing and collecting information about internet censorship. These platforms differ in approach and methodology. At this point however, none of them supports QUIC based protocols, i.e. HTTP/3 or DNS-over-QUIC. ICLab [15] is a censorship observation tool with a measuring approach and structure which is similar to the OONI probing software. The open source project has a vantage point infrastructure, which mainly consists of VPN clients and VPSs, as well as embedded systems. CensoredPlanet [21] is a more recent project that follows a different measurement methodology. Instead of implementing a decentralized approach as ICLab, CensoredPlanet focuses on remote measurement techniques, which can detect connection blocking without controlling either end point. This remote technique exploits TCP/IP side channels by performing a type of reflection attack. The remote measuring approach has the clear advantage of allowing significantly increased coverage. At this point they have around 95,000 vantage points in total. Measurements can be run and replicated without relying on volunteer experiments. It has to be noted, that the technique requires careful ethical considerations. Since the technique exploits side channels, the results are not as specific as traditional measurements. Also, the owners of the examined machines have no knowledge about their participation and could become targeted by the censors.

4 MEASUREMENT FRAMEWORK

The following section describes the measurement framework, the target host list used to create the data set, and the data collection process itself.

4.1 HTTP/3 Extension for OONI Probe

The QUIC censorship measurement software was implemented as part of the *OONI Probe* engine [16], which is written in Go. It consists of a censorship measurement library and a command line interface to conduct multiple networking experiments. Our new QUIC extension can be used in any of the existing OONI network experiments.

To measure website blocking, as done in this work, we configure the existing URLGetter experiment to automatically use the QUIC code. At runtime, we perform a set of preconfigured steps for each entry of the URL test list:

- We parse the URL template to determine next steps.
- We resolve the IP address of the domain name, using either the configured custom resolver or the default system resolver. This step can be replaced by providing a pre-resolved IP address for an input URL, which we have done for the measurements of this work.
- We establish a connection to the host over the configured transport protocol, and try to fetch the resource over HTTP.
- We capture, classify and save any occurring network events or thrown errors during the connection setup, the cryptographic handshake, or the HTTP session.

For the QUIC implementation, we integrated the open source library *quic-go*[1], because it is also written in Go and has been successfully used for QUIC research before [7, 14, 18].

4.2 Vantage Points

The data set was collected by using three different types of clients: Personal Devices of volunteers (PD), Virtual Private Networks (VPNs), and Virtual Private Servers (VPSs).

Personal devices (PD): With the help of OONI volunteers in Iran and India, data was collected on devices from ISPs in AS48147, AS38266 and AS55836. While such data most closely resembles the real conditions in the probed countries, each measurement must be conducted manually by the respective volunteer. Thus, sample size, frequency, and continuity of the data is low. Also, PD measurements invoke risks to the volunteers. Thus, we extended the data set by using remotely controlled VPN and VPS measurements.

Virtual Private Networks (VPNs): Using a VPN, the measurement software can be run on a machine outside the probed network without the need for local volunteers. For measurements in Turkey, Kazakhstan, Russia, and Malaysia, the use of OpenVPN services was considered. However, early measurements showed that the tested VPN servers located in Turkey, Russia, and Malaysia were notably less censored than expected from prior OONI measurements[2] and the Freedom on the Net ranking [5]. This phenomenon is likely caused by one of the following two reasons. First, we observed that most VPN servers are connected to a hosting network, and not to the common local ISP network. Since hosting networks are not used by the general public, it is possible that they are not obligated to implement the same strict censorship policies. The second explanations is that the upstream AS of the VPN is often located in a country with less internet restrictions. In such a scenario, the traffic might never cross a severely censored network in the country of the respective VPN server. To avoid such bias, we did not further conduct measurements in Turkey, Malaysia, and Russia. However, in Kazakhstan, an available VPN server is located in the network of the largest national ISP, KazakhTelecom (AS9198). Since the upstream network of the VPN server's network also belongs to KazakhTelecom, the server can be used as a measurement point for Kazakh internet censorship.

Virtual Private Servers (VPSs): When using a VPS, the application is executed on a virtual machine. This way, various countries can be accessed through OONI and its partners. Some countries, e.g., Costa Rica, Cambodia, and Tunisia, were excluded from our evaluation, because the VPS measurements again showed less interference than expected, similar to our initial VPN measurements. In China, India, and Iran, we used one VPS each in AS45090, AS14061, and AS62442. We expect these VPSs to be affected by the same amount of censorship as traditional measurements on personal devices. We have come to this conclusion after initial measurements which resembled user experience in the respective countries. Because VPS measurements can be taken frequently and continuously, they are crucial to collect a significant number of samples.

4.3 Country-specific Host Lists

To gather meaningful results, the tested websites have to be relevant in the context of internet censorship and also support HTTP/3. To check both criteria, we first created a list of relevant domains, then excluded the ones that did not support QUIC. As mentioned in Section 2, some domain categories were also excluded (e.g., sites with pornographic content). Our starting list of relevant domains includes commonly accessed international domains, websites with controversial topics or country-specific restricted content, and also sites that have been reported as censored in the past. The main source for these domains are the censorship test lists of the Citizen Lab Project[3], often used in the literature [8, 15, 21]. These lists include 1400 mostly English-speaking websites as well as multiple country-specific websites. Additionally, we added to our base list the first 4000 entries of the Tranco Top 1 Million list[4], which ranks the most frequently accessed domains on the web. The so created base list is then filtered by making a QUIC request with cURL[5] and dropping all domains that did not support QUIC. This reduced the final host list significantly. Only about 5% of relevant domains passed. Because some hosts appear to have very unstable QUIC support, we checked all remaining domains one more time in the post-processing step (c.f., Section 4.4).

The composition of each country-specific host list is presented in Figure 2. It has to be noted that the data set contains a significant amount of .com top-level domains. QUIC is currently mainly deployed by large internet companies, e.g., Google, and therefore

[1] https://github.com/lucas-clemente/quic-go
[2] https://explorer.ooni.org/
[3] https://github.com/citizenlab/test-lists
[4] https://tranco-list.eu/
[5] https://curl.se/

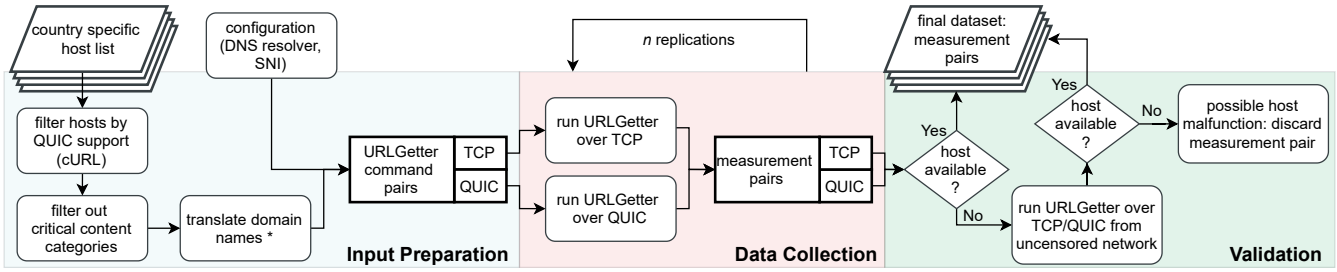

*The DNS step is done via Google DoH from an uncensored network.

Figure 1: Three-step workflow of the conducted censorship measurements

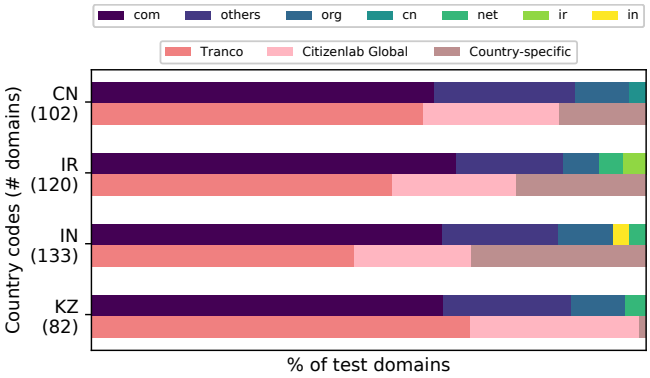

Figure 2: Distribution of top-level domains (first horizontal bar) and sources (second horizontal bar) within each country-specific host list.

predominantly used by globally popular hosts. Due to this bias, we expect to generally see more censorship in countries that largely censor global targets, e.g. China.

4.4 Data Collection Process

As depicted in Figure 1, the final measurement process itself consists of three consecutive phases: input preparation, data collection, and post-processing.

1. Input Preparation: The objective of this step is to obtain a set of *request pairs*. Each pair consists of two HTTP requests to the same target host, one via TCP with TLS and one over QUIC. Such a request pair shares the same configuration parameters, i.e., the content of the TLS extension Server Name Indication (SNI), the pre-resolved IP address of the target host, and the address of a public DoH resolver. The latter DNS configurations ensure that the measurements are not biased by DNS manipulation. The requests are saved as JSON objects and used as input for OONI Probe.

2. Data Collection: All measurements were conducted between January 15th and March 30th 2021 at the vantage points described in Section 4.2. At each VPS vantage point, the entire input list was processed in 8 hours intervals. But due to load variance at the VPSs and temporary server downtime, these intervals shifted sometimes a bit. Table 1 denotes the number of replications and final sample size for each tested AS.

For each request pair, two measurements are performed sequentially, first using TCP, than using QUIC. There is no wait time between the two measurements. A single measurement is done by running the URLGetter experiment on the OONI client with the aforementioned prepared input. The returned report data is then sent to the OONI backend, where it is published via the OONI Explorer API[6].

3. Post Processing & Validation: The post-processing step is necessary, because the QUIC support of some hosts is not stable. Sometimes, random handshake timeout error occur, which are unfortunately not directly distinguishable from timeout errors caused by censorship. To solve this issue, we tested each failed request one more time from an uncensored network before adding the result to the final data set.

If the request fails again, a malfunction at the destination host is assumed and the measurement pair is discarded, including the corresponding QUIC request as well as the TCP request.

5 EVALUATION

This section evaluates the state of QUIC blocking by presenting the measurement results from vantage points in India, China, Iran, and Kazakhstan. We use the HTTPS measurements as a baseline in order to investigate whether the censor blocks both HTTPS and HTTP/3 connections alike. The tangible measurement results are network failures observed at each connection attempt, e.g. *TLS-hs-to*. We use the failure types to derive the most probable failure cause. In case of censorship, the type of the failure can help identify the applied blocking method, e.g. black holing with SNI-based traffic identification.

An overview of our results is shown in Table 1. Notice that the final sample size can be smaller than number of hosts times replications, because invalid requests were filtered out during the validation step.

The very low failure rates of HTTP/3 in AS9198 (Kazakhstan), AS14061, and AS38266 (India) indicate the absence of QUIC censorship at those vantage points. In contrast, the measurements in China, Iran, and India indicate HTTP/3 censorship in the investigated networks. However, we have found that not all domains that are unavailable over TCP, are also unavailable over QUIC. This difference is most prominent is AS45090 (China), where TCP hosts cannot be accessed in 37.3% of cases, whereas only 27.1% of QUIC

[6]https://explorer.ooni.org/

Table 1: Failure rates and error types of connection attempts via HTTPS over TCP and HTTP/3 over QUIC.

Country, (ASN)	Vantage Type, Hosts	Replications, Sample Size*	Failed Attempts						
			TCP					QUIC	
			overall	TCP-hs-to	TLS-hs-to	route-err	conn-reset	**overall**	QUIC-hs-to
China (45090)	VPS, 102	69, 6706	**37.3%**	25.9%	2.7%	-	8.6%	**27.1%**	27.0%
Iran (62442)	VPS, 120	36, 3887	**34.4%**	-	33.4%	-	-	**16.2%**	15.1%
India (55836)	PD, 133	2, 266	**15.0%**	7.5%	-	4.5%	3.0%	**12.0%**	12.0%
India (14061)	VPS, 133	60, 7531	**16.3%**	-	-	-	16.3%	**0.2%**	0.1%
India (38266)	PD, 133	1, 133	**12.8%**	-	-	-	12.8%	-	-
Kazakhstan (9198)	VPN, 82	22, 1764	**3.2%**	-	3.2%	-	-	**1.1%**	1.1%

* final sample size of all replications after validation step filtering (c.f., Figure 1)

Table 2: Decision chart to determine the censor's most likely traffic identification method for a tested domain.

	Response	Additional observation	Conclusion for tested domain	Indication
HTTPS	success	-	no HTTPS blocking	-
	TCP-hs-to, route-err	-	no TLS blocking	IP[1]
	TLS-hs-to, conn-reset	success w/ spoofed SNI	SNI-based TLS blocking, no IP-based blocking	UDP[2]
	TLS-hs-to, conn-reset	failure w/ spoofed SNI	no SNI-based blocking	-
HTTP/3	success	available over HTTPS	no HTTP/3 blocking	-
	success	blocked over HTTPS	HTTP/3 blocking not yet implemented	-
	failure	other HTTP/3 hosts are available in the network	no general UDP/443 blocking in network	UDP[2]
	failure	available over HTTPS	probably blocked as collateral damage	UDP[2]
	QUIC-hs-to	success w/ spoofed SNI	SNI-based QUIC blocking, no IP-based blocking	-
	QUIC-hs-to	failure w/ spoofed SNI	no SNI-based QUIC blocking	IP[1], UDP[2]

[1] Strong indication for IP-based blocking in China and India (c.f., Section 5.1).
[2] Strong indication for UDP endpoint blocking in Iran (c.f., Section 5.2).

hosts are unavailable. Similarly, in the Iranian network AS62442, the failure rate drops from 34.4% to 16.2% when using QUIC instead of TCP. QUIC blocking was also observed in the network of AS55836 in India where 15% of TCP/TLS connections fail and 12% of the HTTP/3 requests. Across all probed networks, the only detected QUIC error type was *QUIC-hs-to*, which suggests the likely use of black holing as an interference method.

Figure 3 breaks down the observed failure types and depicts how the connection response of each tested host changes when using QUIC (right-hand side) instead of TCP. Subsections 5.1 and 5.2 further explain this Figure, and provide a more detailed analysis of the results.

5.1 IP-based Blocking

IP blocking is one of the most common censorship methods. Since traffic is blocked depending on the destination IP, it affects QUIC and TCP traffic alike. From our measurements, we derive the presence of IP-based blocking in AS45090 (China) and AS55836 (India). The corresponding pieces of evidence are listed in Table 2 and will be explained in the following.

Due to the larger sample size and temporal consistency, this section focuses especially on the AS45090 (China) data set. There, three different types of network errors were measured: *TCP-hs-to*, *TLS-hs-to*, and *conn-reset* errors (c.f., Table 1). All hosts, that raised an HTTPS connection reset error are still available via HTTP/3 over QUIC. Similar, in the case of TLS handshake errors over HTTPS, the corresponding HTTP/3 attempt nearly always succeeds. These findings indicate that HTTP/3 over QUIC traffic is less censored than traditional HTTPS via TCP, in AS45090. However, if the HTTPS request times out during the TCP handshake (*TCP-hs-to*), an HTTP/3 request also fails before the QUIC handshake completes.

In such scenarios in AS45090, it is most likely that the destination IP addresses is blocked: Because neither handshake is completed, TLS-based censorship techniques, such as SNI-based blocking can be ruled out. General TCP and UDP port blocking on 443 can also be disregarded, because other HTTPS and HTTP/3 requests during the same measurement round succeeded. This leaves IP blocking as the most probable explanation. Unfortunately the usage of HTTP/3 over QUIC can not overcome this type of censorship, because the interference already happens on the underlying IP layer. Since the censor does not exclusively apply IP endpoint blocking, hosts that are targeted by a different form of HTTPS censorship are still available over QUIC.

IP blocking was also observed in the network of an Indian ISP located in AS55836. As shown in Figure 3b, for every TCP connection error associated with IP-blocking, (*TCP-hs-to* and *route-err*), the corresponding QUIC measurement also fails. This indicates that, like in AS45090, the applied IP-blocking affects QUIC in the same way as TCP but does not target the protocol directly.

5.2 UDP Endpoint Blocking

In Iran (AS62442), most HTTPS errors occur due to *TLS-hs-to's*, i.e., timeouts after the establishment of a TCP connection. This is not a case of IP blocking. Instead, the TLS handshake timeout rather indicates that a TLS-blocking method is active, such as commonly used SNI-filtering. As depicted in Figure 3c, a third of the unsuccessful HTTPS attempts also fail if HTTP/3 is used instead, returning a timeout during the QUIC handshake (*QUIC-hs-to*).

To test our hypothesis, that the SNI field is used as host identification for TLS-blocking in Iran, a subset of the host list was additionally probed with the SNI field in the *ClientHello* set to *example.org*. A similar approach is used in [2]. As displayed in Table 3, 83% of attempts to usually blocked TCP/TLS hosts succeed when applying the SNI spoof. This behavior suggests that Iranian censors utilize SNI keyword filtering to block TCP/TLS connections and

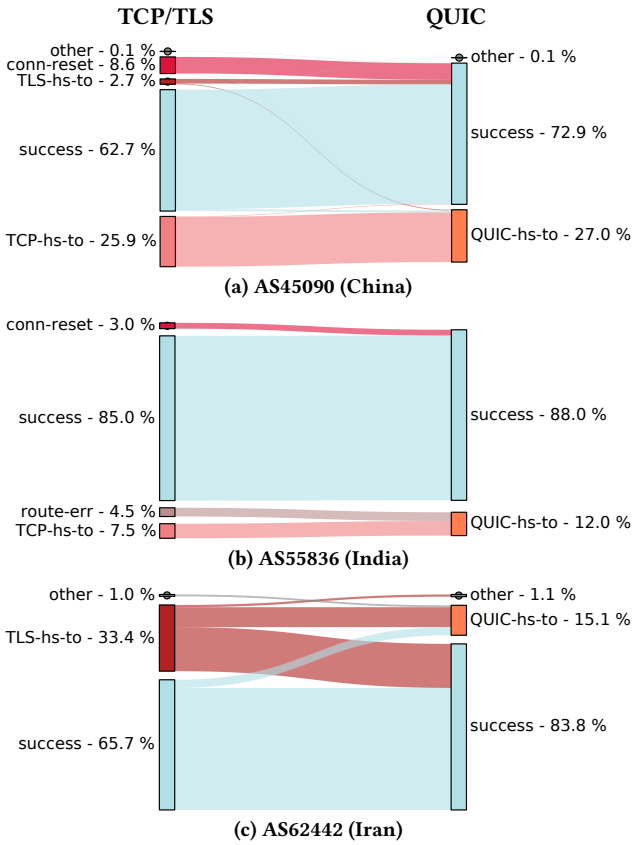

Figure 3: Distribution of network error types for TCP/TLS (left) and QUIC (right) measurements. The horizontal flows (left to right) indicate the response change, when using QUIC instead of TCP/TLS.

rules out general IP-endpoint blocking. However, using the spoofed SNI does not affect the availability of the subset hosts when using HTTP/3.

In comparison to the other probed networks, the percentage of pairs with a successful TCP/TLS attempt and a failed QUIC attempt, is more pronounced, totaling 4.11% of all pairs. The additional availability test in the post-processing step showed that the affected hosts are still available via QUIC from an uncensored AS, ruling out any server-side UDP fire walling. Since the majority of website requests is still made with traditional HTTPS, it seems unreasonable that a censor would intentionally only block the lesser used protocol, and we therefore assume that these hosts are unavailable as collateral damage caused by IP address filtering.

Thus, we believe that censors have deployed middle box software, which applies IP address filtering only to UDP traffic. As described above and listed in Table 2, we derive this conclusion from the elimination of IP-based blocking and SNI blocking, and from the observed collateral damage in regards to censored QUIC connections. Future work has to prove, if this filter specifically targets HTTP/3 traffic, i.e. UDP traffic on port 443, or UDP traffic to these IPs in general.

Table 3: SNI-based TLS blocking and SNI spoofing measurements in Iran.

ASN	transport	sample size	failure rate	
			real SNI	spoofed SNI (*example.org*)
62442	TCP	353	60.1% (212)	10.2% (36)
(Iran)	QUIC	353	20.1% (71)	20.1% (71)
48147	TCP	40	60.0% (24)	10.0% (4)
(Iran)	QUIC	40	20.0% (8)	20.0% (8)

6 CONCLUSION & FUTURE WORK

The presented findings summarize the current state of QUIC censorship in several critical Autonomous Systems. Our measurements reveal that HTTP/3 requests are less frequently blocked than traditional HTTPS requests, or sometimes even not blocked at all. This can be explained by the fact QUIC is a rather young protocol with still a small share in overall internet traffic volume. Yet, with its growing significance, the efforts to better block QUIC will rise. As observed with the outright blocking of Encrypted-SNI in China[7], it is also possible that QUIC could be generally blocked by censors.

Two identification methodologies were observed in the probed networks, aiming indirectly and directly at QUIC traffic. The blocklisting of IP addresses is still an ongoing issue, as seen in the probed Chinese network, and in one of the probed networks in India. While such IP blocking prevents HTTP/3 requests to blocklisted hosts, it also affects all other IP-based protocols. In the two Iranian networks, we detected UDP endpoint blocking used against HTTP/3 connections. This method differs from the applied TLS-blocking method which filters HTTPS traffic based on the SNI. The only observed interference method used to intercept QUIC connections is black holing to interrupt the handshake.

Since censorship methods dynamically change and censors adapt to the emergence of new network technologies, measurements can only reflect the censorship situation at a certain point in time. Currently, QUIC is not fully deployment internationally, as discussed in our input selection, c.f., Section 4.3. Thus, this work only presents a fixed snapshot of an early stage of QUIC censorship. The study should be repeated in near future to highlight the development.

Beyond the collected data, this work provides a measurement tool to long-term monitor HTTP/3 over QUIC blocking around the world. Future measurements should not only monitor the use of established censorship methodologies applied to QUIC, but also stay alert to detect new methods tailored to QUIC or TLS traffic and identify the use of statistical flow classification.

ACKNOWLEDGMENTS

We want to acknowledge the work and commitment of the OONI community, which helped in conducting the measurements. In particular, we thank Gurshabad Grover and Divyank Katira from CIS India for their efforts. We also want to acknowledge the organization GreatFire.org, which enabled us to perform measurements in China. We thank the shepherd Mirja Kuehlewind and the anonymous reviewers for their helpful feedback.

[7] https://gfw.report/blog/gfw_esni_blocking/en/

REFERENCES

[1] Simurgh Aryan, Homa Aryan, and J. Alex Halderman. 2013. Internet Censorship in Iran: A First Look. In *3rd USENIX Workshop on Free and Open Communications on the Internet (FOCI 13)*.

[2] Simone Basso, Gurshabad Grover, and Kushagra Singh. 2020. Investigating TLS blocking in India. (2020). Available: https://ooni.org/post/2020-tls-blocking-india/.

[3] Mike Bishop. 2021. *Hypertext Transfer Protocol Version 3 (HTTP/3).* Internet-Draft draft-ietf-quic-http-34. IETF. Available: https://datatracker.ietf.org/doc/html/draft-ietf-quic-http-34.

[4] Kevin Bock, Yair Fax, Kyle Reese, Jasraj Singh, and Dave Levin. 2020. Detecting and Evading Censorship-in-Depth: A Case Study of Iran's Protocol Whitelister. In *10th USENIX Workshop on Free and Open Communications on the Internet (FOCI 20)*.

[5] Noah Buyon, Cathryn Grothe, Amy Slipowitz, and Kian Vesteinsson. 2020. Freedom on the Net. (2020). Available: https://freedomhouse.org/sites/default/files/2020-10/10122020_FOTN2020_Complete_Report_FINAL.pdf.

[6] Zimo Chai, Amirhossein Ghafari, and Amir Houmansadr. 2019. On the Importance of Encrypted-SNI (ESNI) to Censorship Circumvention. In *9th USENIX Workshop on Free and Open Communications on the Internet (FOCI 19)*.

[7] Quentin De Coninck and Olivier Bonaventure. 2017. Multipath QUIC: Design and Evaluation. In *Proceedings of the 13th International Conference on emerging Networking EXperiments and Technologies (CoNEXT '17)*. 160–166.

[8] Arturo Filastò and Jacob Appelbaum. 2012. OONI: Open Observatory of Network Interference. In *2nd USENIX Workshop on Free and Open Communications on the Internet (FOCI 12)*.

[9] Joseph Hall, Ben Jones, Michael Aaron, Amelia Andersdotter, Stan Adams, and Nick Feamster. 2020. *A Survey of Worldwide Censorship Techniques.* Internet-Draft draft-irtf-pearg-censorship-02. IETF. Available: https://tools.ietf.org/html/draft-irtf-pearg-censorship-02. Expired September 10, 2020 .

[10] Kyle Hogan. 2020. *Security analysis of Tor over QUIC.* PhD Thesis. Massachusetts Institute of Technology. Available: https://dspace.mit.edu/handle/1721.1/128590.

[11] Jana Iyengar and Martin Thomson. 2021. *QUIC: A UDP-Based Multiplexed and Secure Transport.* Internet-Draft draft-ietf-quic-transport-34. IETF. Available: https://datatracker.ietf.org/doc/html/draft-ietf-quic-transport-34.

[12] Adam Langley, Alistair Riddoch, Alyssa Wilk, Antonio Vicente, Charles Krasic, Dan Zhang, Fan Yang, Fedor Kouranov, Ian Swett, Janardhan Iyengar, Jeff Bailey, Jeremy Dorfman, Jim Roskind, Joanna Kulik, Patrik Westin, Raman Tenneti, Robbie Shade, Ryan Hamilton, Victor Vasiliev, Wan-Teh Chang, and Zhongyi Shi. 2017. The QUIC Transport Protocol: Design and Internet-Scale Deployment. In *Proceedings of the Conference of the ACM Special Interest Group on Data Communication (SIGCOMM '17)*. 183–196.

[13] B. Martini and N. ten Oever. 2018. *QUIC Human Rights Review.* Internet-Draft draft-martini-hrpc-quichr-00. IETF. Available: https://tools.ietf.org/id/draft-martini-hrpc-quichr-00.html. Expired April 25, 2019.

[14] François Michel, Quentin De Coninck, and Olivier Bonaventure. 2019. QUIC-FEC: Bringing the benefits of Forward Erasure Correction to QUIC. In *Proceedings of the IFIP Networking 2019 Conference*. 1–9.

[15] A. A. Niaki, S. Cho, Z. Weinberg, N. P. Hoang, A. Razaghpanah, N. Christin, and P. Gill. 2020. ICLab: A Global, Longitudinal Internet Censorship Measurement Platform. In *2020 IEEE Symposium on Security and Privacy (SP)*. 135–151.

[16] Open Observatory of Network Interference. 2021. ooni/probe-cli. (2021). Available: https://github.com/ooni/probe-cli.

[17] Reethika Ramesh, Ram Sundara Raman, Matthew Bernhard, Victor Ongkowijaya, Leonid Evdokimov, Anne Edmundson, Steven Sprecher, Muhammad Ikram, and Roya Ensafi. 2020. Decentralized control: A case study of russia. In *Network and Distributed Systems Security Symposium (NDSS)*.

[18] Jan Rüth, Ingmar Poese, Christoph Dietzel, and Oliver Hohlfeld. 2018. A First Look at QUIC in the Wild. In *Passive and Active Measurement*, Vol. 10771. 255–268.

[19] Kushagra Singh, Gurshabad Grover, and Varun Bansal. 2020. How India Censors the Web. In *12th ACM Conference on Web Science (WebSci '20)*. 21–28.

[20] Ram Sundara Raman, Leonid Evdokimov, Eric Wurstrow, J. Alex Halderman, and Roya Ensafi. 2020. Investigating Large Scale HTTPS Interception in Kazakhstan. In *Proceedings of the ACM Internet Measurement Conference (IMC '20)*. 125–132.

[21] Ram Sundara Raman, Prerana Shenoy, Katharina Kohls, and Roya Ensafi. 2020. Censored Planet: An Internet-wide, Longitudinal Censorship Observatory. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security (CCS '20)*. 49–66.

[22] Ram Sundara Raman, Adrian Stoll, Jakub Dalek, Reethika Ramesh, Will Scott, and Roya Ensafi. 2020. Measuring the Deployment of Network Censorship Filters at Global Scale. In *Proceedings of the Network and Distributed System Security Symposium, NDSS 2020*.

[23] Pengwei Zhan, Liming Wang, and Yi Tang. 2021. Website Fingerprinting on Early QUIC Traffic. *arXiv:2101.11871 [cs]* (2021).

NOTICE

Ethical considerations are discussed in Section 2.