IRBlock: A Large-Scale Measurement Study of the Great Firewall of Iran

Jonas Tai and Karthik Nishanth Sengottuvelavan, *University of British Columbia;* Peter Whiting, *University of Waterloo;* Nguyen Phong Hoang, *University of British Columbia*

https://www.usenix.org/conference/usenixsecurity25/presentation/tai

This paper is included in the Proceedings of the 34th USENIX Security Symposium.

August 13–15, 2025 • Seattle, WA, USA

978-1-939133-52-6

IRBlock: A Large-Scale Measurement Study of the Great Firewall of Iran

Jonas Tai* Karthik Nishanth Sengottuvelavan* Peter Whiting† Nguyen Phong Hoang*

*University of British Columbia †University of Waterloo

Abstract

The Great Firewall of Iran (GFI) has evolved significantly over the past decade, constantly adding sophisticated blocking techniques. Prior research into Iran's Internet censorship, however, has primarily been one-off studies, leaving significant gaps in understanding the breadth and evolution of its filtering strategies. Exploiting the bidirectional blocking behaviors of the GFI and its own injection mechanisms as a side-channel to determine traffic disruption, we developed `IRBlock`, a novel large-scale, multi-protocol measurement system designed to measure DNS, HTTP, and UDP-based censorship across Iran, enabling continuous monitoring and in-depth exploration of the GFI's blocking behavior.

Over a period of 2.5 months, `IRBlock` has periodically measured the entire Iran's IP address space and tested the blocking status of over 500M apex domains, uncovering new insights into the GFI's censorship practices of different core network protocols. Notably, `IRBlock` identified 6.8M IPs subjected to DNS poisoning and HTTP blockpage injection, and 5.4M IPs subjected to UDP-based traffic disruption. We also analyzed the censored domains found by `IRBlock` and discovered over censored 6M FQDNs and 3.3M apex domains. Via reverse engineering of the GFI's blocking rules, we found many domains are inadvertently overblocked due to blanket blocking policies of entire TLDs (e.g., `.il`), resulting in large collateral damage to innocuous websites. We also find that the GFI's blocking strategies show many similarities to those observed for the Great Firewall of China.

Our study represents the most comprehensive view of Iran's Internet censorship to date. Leveraging `IRBlock`'s data, we shed light on the GFI's evolving filtering strategies and the challenges faced by circumvention tools. We discuss the implications of our findings on existing censorship measurement and circumvention efforts. We hope that the insights gained from our study can inform not only the research community but also policymakers and activists working to promote digital freedom in Iran and beyond. All data collected by `IRBlock` will be made publicly available to facilitate further research on nation-state censorship and Internet freedom advocacy.

1 Introduction

As reported in the 2024 Freedom on the Net by the Freedom House, Iran is among the world's most aggressive Internet censors [39]. The country ranks third in the world for lowest Internet freedom score, behind only China with the notorious Great Firewall (GFW) [40] and Myanmar, a Southeast Asian country with a history of military rule and human rights abuses [32, 57]. Iran's censorship apparatus is complex, employing multi-layered filtering mechanisms, including DNS, HTTP, HTTPS filtering [23], and protocol-based blocking [29] (§2.1). There have also been reports of traffic throttling [21]. And over the past decade, the Great Firewall of Iran (GFI) [1] has evolved significantly, constantly adding new and sophisticated blocking techniques as new protocols and applications emerge.

As nation-state censorship systems have gotten more sophisticated, impacting the freedom of information and expression online, especially in repressive regions of the world, the Internet freedom and censorship research community has made significant strides in developing new methods to measure and better understand these systems. The Open Observatory of Network Interference (OONI) [37], ICLab [51], and Censored Planet [63] are a few global measurement platforms that have been instrumental in monitoring Internet censorship around the world. Some other recent studies, including Xue et al. [69], Ramesh et al. [64], Raman et al. [61], GFWatch [44], GFWeb [43], and TMC [54], have made efforts in understanding nation-state censorship systems in countries like Russia, China, Kazakhstan, and Turkmenistan. However, research on Iran's Great Firewall (GFI) remains sparse due to many measurement challenges (§2.2), with only a handful of studies focusing on specific protocols or short timeframes [21, 23, 29].

As we will discuss in more detail in §2.2, one of the key challenges that has limited the scope of existing studies is

[1] We coin the term the Great Firewall of Iran (GFI) to refer to Iran's national censorship system, analogous to the Great Firewall of China (GFW). Throughout this paper, we will show how the GFI employs several blocking strategies with similarities to those of the GFW.

the need for vantage points inside the target country. Renting virtual private servers (VPS) or using commercial VPNs as vantage points is not feasible in countries like Iran, where it is extremely difficult to obtain a server for experiments due to US sanctions [55] and the country's strict Internet regulations. As a result, a systematic understanding of the GFI's blocking behavior and evolution is lacking, with most of the existing knowledge coming from studies that are limited in scope and duration [21, 23, 29] and from measurement platforms like OONI [37] and Censored Planet [63] which rely on volunteers and public servers, leaving significant gaps in our understanding of the GFI.

To complement existing efforts and address these gaps, we developed `IRBlock`, a large-scale, multi-protocol measurement system designed to detect DNS poisoning, HTTP blockpage injection, and UDP-based filtering in Iran. Importantly, by exploiting the bidirectional blocking of the GFI (§4) and its own injection mechanisms as a side-channel to determine traffic disruption, `IRBlock` is capable of measuring the entire Iranian IP space, which consists of millions of IPs and hundreds of millions of domains on a regular basis. This allows us to provide a more comprehensive view of the GFI's censorship practices, including the identification of censored IPs and domains, the discovery of new blocking mechanisms, and the analysis of the GFI's blocking rules and behavior.

Over a monitoring period of 2.5 months reported in this paper, `IRBlock` has periodically measured the entire Iran's IP address space of over 11M IPs, identifying over 6.8M IPs that experience DNS poisoning and HTTP blockpage injection, and over 5.4M IP addresses that experienced UDP-based traffic disruption (§5.1). By analyzing the blocking status of over 700M fully-qualified domains (FQDNs), we collect the largest blocklist of the GFI ever discovered to date, revealing over 6M banned FQDNs, originating from over 3.3M censored apex domains (§5.3).

Data collected by `IRBlock` also reveals blanket blocking of several top-level domains (TLDs), such as `.*\.il$` and `.*\.porn$` (§5.3). We also discover three different injectors with different blocking signatures, and blocking behaviors of the GFI that parallel strategies observed in the GFW, such as the reflection of TTL values in middlebox responses (§5.4).

All data collected by `IRBlock` will be made publicly available to foster further research on nation-state censorship and support Internet freedom advocacy. Our ongoing efforts aim to continuously expand our dataset, providing regular updates on the GFI's censorship practices. By providing a more comprehensive view of the GFI's blocking behavior, we hope to inform not only the research community but also policymakers and activists working to promote digital freedom in Iran and beyond.

2 Background and Motivation

Next, we detail several core blocking mechanisms of the Great Firewall of Iran (GFI), including DNS poisoning, HTTP blockpage injection, and HTTPS connection tear-down, and UDP traffic dropping. We then discuss the challenges of measuring Iran censorship faced by existing studies and how that motivated us to design `IRBlock` to overcome these challenges, complementing existing efforts to provide a more comprehensive view of the GFI's censorship practices.

2.1 The GFI's Censorship Mechanisms

The GFI employs multiple blocking mechanisms to censor Internet traffic. In 2013, Aryan et al. [23] were among the first groups to study the GFI's inner working, detecting DNS poisoning and HTTP blocking, as well as connection throttling. Since then, other studies have observed TCP-based protocol blocking [29], and HTTP(S) blocking as part of global censorship studies [30, 62]. Based on these prior works and recent anecdotal reports [1, 5], we summarize the GFI's censorship mechanisms in Figure 1 to illustrate the four primary blocking mechanisms from a client's perspective.

DNS Poisoning. A common, lightweight censorship technique is the injection of fake DNS responses. The original DNS protocol [50], running on top of UDP, is an unencrypted, stateless protocol. Making use of this design, a censoring middlebox can inspect DNS packets and inject falsified responses. As illustrated in Figure 1(a), a DNS query carrying a censored domain (e.g., `hrw.org`) is intercepted and dropped since the GFI functions as an in-path filtering middlebox. Thus, the DNS query never reaches the actual DNS resolver, unlike the GFW that operates as an on-path filter and still lets the censored query and the actual DNS response pass through [22, 44]. Upon detecting the censored query, the GFI then sends a forged DNS reply to the user's device, containing a private IP address in the form of `10.10.x.x`, but routable inside Iran [20], redirecting them to a local blockpage. As we will show in §5, there are three different DNS injectors whose forged replies contain different fake IP addresses.

This DNS poisoning technique is a common censorship strategy used by many countries due to its simplicity and effectiveness, including China [22, 44] and Turkmenistan [54]. However, as we will show in §5, the existence of multiple DNS injectors in Iran, each with its own unique blocking signatures (e.g., injected forged IPs and TTL values), is similar to blocking strategies used by the GFW [22, 44].

HTTP Blockpage Injection. As a multi-layered censorship system, the GFI also employs HTTP-based filtering to block access to censored websites in case the DNS poisoning fails or is circumvented. As shown in Figure 1(b), the GFI tracks the initial TCP handshake and inspects the first HTTP packet, which includes the unencrypted domain name in the `HTTP`

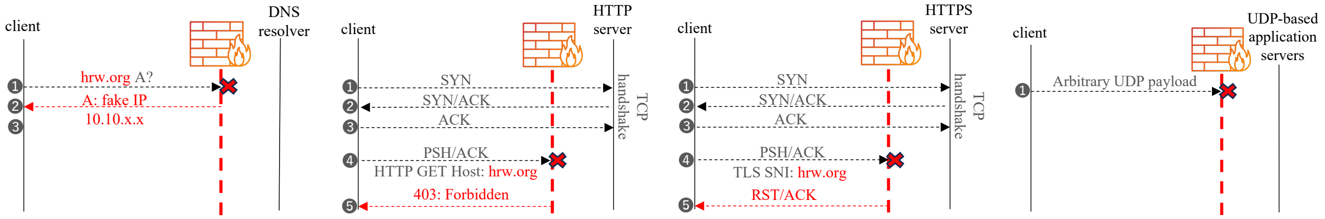

(a) DNS filter

(b) HTTP filter

(c) HTTPS filter

(d) UDP filter

Figure 1: The DNS, HTTP, and HTTPS filtering mechanisms of the GFI. SYN, ACK, PSH, and RST denote TCP synchronization, acknowledgement, push, and reset flags. A packet with the RST flag set is meant to terminate a TCP connection.

GET request header. If the domain is found in the HTTP filter's blocklist, the GFI injects an HTTP 403 - Forbidden response back to the source. The HTTP blockpage injection is a more complicated mechanism than DNS poisoning, as it requires deep packet inspection and a stateful firewall to track the initial TCP handshake. This is different from other systems like the GFW and Turkmenistan's censorship system, which use RST packets to terminate the connection [43, 54].

Similar to DNS poisoning, the HTTP blockpage injection also has multiple injectors, each with its unique blocking signatures which we will analyze in more detail in §5.

HTTPS Connection Disruption. Next, if the user (1) attempts to access a censored website over HTTPS by typing https:// in the URL bar instead of http:// or their browser automatically enforces HTTPS, or (2) could circumvent the HTTP blockpage injection [41], the GFI next employs a more sophisticated HTTPS connection disruption mechanism. As shown in Figure 1(c), the GFI inspects the unencrypted Server Name Indication (SNI) field of the TLS client hello after the TCP handshake. If the domain is found in the HTTPS filter's blocklist, the GFI sends an injected RST+ACK packet to terminate the connection.

The observed HTTP and HTTPS censorship mechanisms require deep packet inspections and stateful tracking of the initial TCP handshake, making them more expensive than DNS censorship. The stateful nature of the GFI for both HTTP and HTTPS filtering as we observed in our experiments is a significant departure from past studies, which reported stateless HTTP and HTTPS blocking [30]. Moreover, we also found that both HTTP and HTTPS filters have been now activated on all TCP ports, differing from prior reports where these filter were limited to standard TCP ports [29, 30].

UDP Traffic Dropping. Recent anecdotes [1, 5] suggest that the GFI also block UDP traffic, which has become increasingly common in modern protocols like QUIC and tunneling protocols (e.g., OpenVPN and WireGuard) often used to circumvent censorship. UDP's connectionless and stateless nature, while enabling efficient and low-latency communication, also presents unique challenges for censorship systems. These characteristics necessitate distinct blocking strategies.

UDP traffic is particularly susceptible to techniques such as packet dropping, where packets are discarded without acknowledgment, or targeted blackholing (null routing) of specific IP addresses associated with suspected circumvention tools. As shown in Figure 1(d), the GFI may drop UDP packets it deems as circumvention traffic in a non-discriminatory manner. This blocking extends to protocols like QUIC, whose encrypted-by-default nature restricts the visibility of payloads, making it harder for censors to inspect the traffic [36]. This lack of visibility into the payload increases the reliance on metadata analysis (e.g., packet size, port number, or traffic patterns) to identify and block specific UDP traffic.

Protocols like QUIC are particularly targeted due to their increasing adoption as the transport layer for HTTP/3 [48]. QUIC reduces connection establishment latency by bundling the cryptographic handshake (TLS ClientHello) into the first packet, encrypting much of its content using initial keys defined in RFC 9000 [45]. This encryption means the Server Name Indication (SNI) is not exposed for easy inspection by middleboxes, unlike traditional TLS over TCP. Consequently, decrypting QUIC packets to identify censored domains requires significant computational overhead, making simple traffic dropping a more resource-efficient censorship strategy. Moreover, the recent anecdotes [1, 5] also suggest that the GFI not only targets QUIC in particular, but also employs broader strategies against UDP traffic to hinder circumvention. Unfortunately, these anecdotal reports lack a comprehensive analysis of the GFI's UDP censorship, providing only a limited understanding of the GFI's UDP blocking behavior from a few vantage points of in-country users.

These observations together indicate that the GFI has evolved significantly over the past decade, constantly adding new and sophisticated blocking techniques, necessitating a more in-depth study to understand its new capabilities.

2.2 Measurement Challenges

Studying the Great Firewall of Iran (GFI) poses unique challenges due to the complexity of its censorship mechanisms, the scarcity of in-country resources, and the variability of its filtering policies across regions and ISPs. Below, we detail the main obstacles faced in measuring the GFI and how they have motivated the design of `IRBlock`.

Limitations of In-Country Access. One of the primary hurdles in studying the GFI is the limited availability of in-country vantage points that can be used for large-scale, longitudinal measurements. The Open Observatory of Network Interference (OONI) [37] and Censored Planet [63] rely on volunteers and public servers to conduct censorship measurements. While these platforms have been instrumental in monitoring Internet censorship around the world, it is challenging to recruit local volunteers that can help with large-scale, long-term measurements due to the risks posed by the country's strict surveillance apparatus. Similarly, sending a large amount of probes to alive public servers may raise potential ethical concerns. Obtaining commercial VPN services, an approach used by ICLab [51], is also not feasible in Iran due to the country's strict Internet regulations and US sanctions [55]. As a result, many aspects of the GFI's censorship practices remain unexplored.

Inconsistent Filtering Policies. Prior studies and anecdotes have reported that the GFI's censorship policies can vary across regions and ISPs [29], complicating efforts to generalize findings and predict behavior, as measurements conducted in one region or through one ISP may not represent the broader censorship landscape. In fact, we will later show in §5.1 that the GFI's blocking behavior can differ significantly across different Autonomous Systems (ASes). Moreover, as shown in §2.1 above, the GFI frequently updates its filtering mechanisms, introducing additional layers of uncertainty for long-term studies.

These challenges have motivated us to design a system that can measure the GFI's censorship practices more comprehensively, without using in-country vantage points. In the next section, we will introduce `IRBlock`, a novel large-scale, multi-protocol measurement system designed to detect DNS poisoning, HTTP blockpage injection, and UDP-based filtering across Iran, enabling unprecedented monitoring and in-depth exploration of the GFI's blocking behavior.

3 The GFI's Unique Blocking Behaviors and Measurement Opportunities

Considering the unique challenges associated with measuring the Great Firewall of Iran (GFI), we designed `IRBlock` with the following two key objectives: (1) to conduct large-scale measurements of the GFI's censorship mechanisms, specifically focusing on core network protocols to provide a more comprehensive view of the GFI's blocking practices, enabling the discovery of as many censored domains and IPs as possible to provide timely insights to the public and research community; and (2) to ensure that our measurement activities do not pose any risks to volunteers or saturate public servers. To overcome the challenges discussed in §2.2 and achieve these goals, we design `IRBlock` to operate entirely from outside Iran, leveraging several unique blocking behaviors of the GFI to measure the entire Iran's IP space, which consists of millions of IPs, and test hundreds of millions of domains on a regular basis, enabling large-scale data collection and continuous monitoring of the GFI's censorship practices without relying on users inside Iran.

3.1 Bidirectional Interference

Recent advances in censorship measurement techniques have demonstrated the potential of exploiting the bidirectional nature of many censorship systems to gain insights into their mechanisms [53], such as the Great Firewall of China and Turkmenistan's censorship system [43, 44, 54]. A censorship system is considered bidirectional when it does not differentiate between incoming and outgoing traffic, thus taking interference actions on both directions regardless of the traffic's origin. Such bidirectional behavior opens up new opportunities for large-scale measurements of nation-state censorship systems, as it allows measurement tools run outside the censored country to observe and infer censorship events without the need for in-country vantage points.

Similar to China and Turkmenistan, the GFI exhibits bidirectional behavior in all foundational Web protocols that we tested, including DNS, HTTP, and HTTPS. In other words, even if the client in Figures 1(a), 1(b), and 1(c) is located outside Iran, the GFI still injects a fake response to the DNS query, a HTTP blockpage to the HTTP request, and a RST+ACK packet to the HTTPS request if the payload contains a censored domain. This is the very behavior that enables Censored Planet [63] to conduct censorship measurement in many countries where there are public DNS resolvers and open HTTP(S) servers.

We also leverage this bidirectional behavior to detect censored domains and IPs by sending probes from outside Iran and inferring censorship by analyzing the responses. Nonetheless, our outside-in measurement approaches are different from prior efforts such as Satellite [65], Iris [59], Quack [68], and Hyper-Quack [67], as we do not rely on responsive open servers for data collection.

3.2 TCP Non-Compliance

While `IRBlock` leverages the bidirectional design of the GFI to measure censorship, similar to some existing systems, we exploit the GFI's non-compliance with standard TCP behavior to eliminate the need for responsive IPs in our measurements. Specifically, our initial investigation revealed that the GFI's HTTP filter does not require a full TCP handshake to trigger filtering, allowing us to send a SYN packet followed by a PSH/ACK packet containing an HTTP GET request for a censored domain. As long as the PSH/ACK packet's sequence

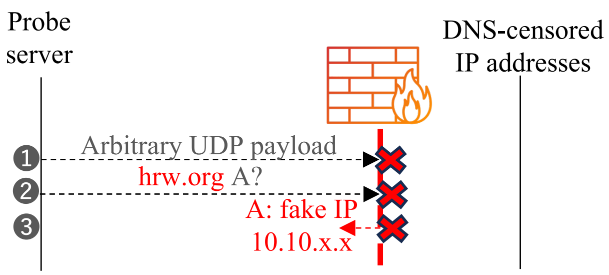

Figure 2: Triggering UDP traffic dropping and detection using the GFI's DNS injection as a side-channel.

number matches the SYN packet's sequence number plus one, the GFI processes the connection as if it were established and injects a HTTP blockpage with a status code of 403 (Forbidden). In other words, a probe machine from outside Iran can trigger HTTP censorship (Figure 1(b)) without receiving a SYN/ACK (packet ❷, which is normally required by prior techniques) or sending an ACK (packet ❸) to complete the TCP handshake, as long as packet ❶ is sent and the PSH/ACK (packet ❹) contains a censored domain and has the correct sequence number.

This TCP non-compliance design is similar to the GFW [43] and Turkmenistan's censorship system [54], most likely to make the censoring middleboxes more resistant to packet loss and dynamic routing. A similar technique was used to measure Internet censorship in China [43] and Turkmenistan [54], and was initially noticed by Bock et al. [28] in the context of amplification attacks. Nevertheless, as of this writing, this TCP non-compliance behavior of the GFI has changed from past studies, where it was previously reported to operate in a stateless manner [30], but has now become stateful (i.e., requiring the initial SYN packet and the matching of sequence numbers to be triggered).

Unlike the case of the GFW [43] and Turkmenistan's censorship system [54], the GFI does not exhibit this similar behavior consistently for HTTPS censorship. While we can trigger HTTPS censorship from outside the country (Figure 1(c)) as long as the other end responds with a SYN/ACK packet (❷), we cannot reliably trigger the RST/ACK packet (❸) that indicates censorship without relying on open public servers. While prior studies have used this technique to measure HTTPS censorship, due to the scale of our measurements to study the entire Iranian IP space and test as many domains as possible, it is ethically unsound to rely on public servers for our measurements. Therefore, we opt to not conduct large-scale measurements of HTTPS censorship in this study, and only leverage the TCP non-compliance of the GFI for measuring HTTP censorship.

3.3 DNS Injection as a Side-Channel for Detecting UDP Traffic Disruption

Recent anecdotal reports [23, 29] have noted that the GFI also employs UDP-based censorship mechanisms, which are challenging to measure without controlling both ends of the testing connection. In fact, most anecdotal reports on UDP traffic dropping are from on-the-ground Iranian users who have machines inside and virtual-private servers outside the country to check whether their UDP packets are dropped. While these efforts provide valuable preliminary insights into the GFI's UDP dropping practices, they are limited in scope and provide inconsistent views due to the small number of vantage points and the limited number of tests conducted. Moreover, some users even admit the potential risks of running such tests from inside Iran due to the country's strict Internet regulations.

Based on these community reports and the way the GFI handles DNS injection, we hypothesize that we can use DNS injection as a side-channel to detect UDP-based censorship without requiring a responsive IP. As shown in Figure 2, when probing a censored IP address with an arbitrary UDP packet, and subsequently sending a DNS query containing a known censored domain, we observe that the GFI drops all future UDP traffic that shares the same tuple of (source IP, source port, destination IP, destination port). This results in no DNS injection observed, which usually occurs when sending a DNS query for a known censored domain to the same IP address.

This could happen because the GFI drops subsequent UDP traffic from the same UDP stream, thus preventing the DNS query from reaching the GFI's DNS filter, or the DNS query does reach the GFI, but the GFI's DNS filter (1) does not respond with an injection due to the ongoing dropping of subsequent UDP traffic or (2) responds with a DNS injection, but the response is dropped by the UDP-dropping module. In either case, the absence of a DNS injection response serves as a reliable indicator of UDP traffic dropping. This observation opens up a new avenue for us to design a novel measurement technique to detect UDP traffic dropping at scale without requiring in-country servers by leveraging the GFI's own DNS injections as a side-channel.

4 *IRBlock*'s Measurement Pipeline

IRBlock's measurement pipeline consists of two main network scanning phases. The initial scans of the entire Iran's IP space are to identify censored IP addresses subjected to DNS and HTTP censorship. Censored IP addresses are then probed to identify domains blocked by the DNS and HTTP filters and detect UDP traffic dropping. In this section, we describe *IRBlock*'s measurement pipeline as shown in Figure 3.

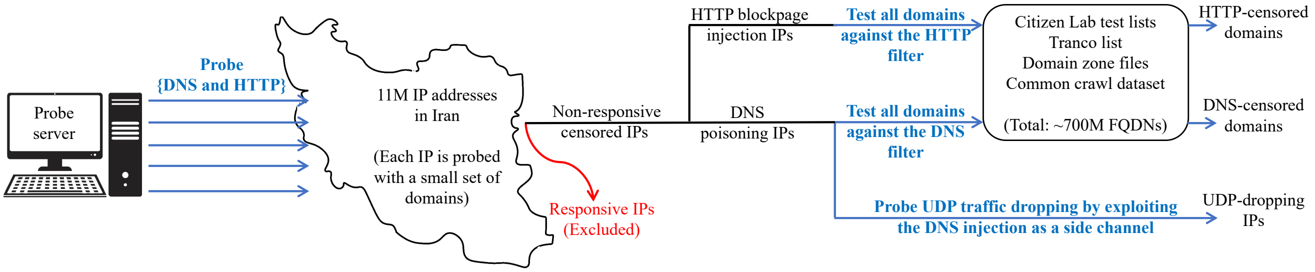
Probe server
Probe
{DNS and HTTP}
11M IP addresses in Iran
(Each IP is probed with a small set of domains)
Non-responsive censored IPs
Responsive IPs (Excluded)
HTTP blockpage injection IPs
DNS poisoning IPs
Test all domains against the HTTP filter
Test all domains against the DNS filter
Citizen Lab test lists
Tranco list
Domain zone files
Common crawl dataset
(Total: ~700M FQDNs)
HTTP-censored domains
DNS-censored domains
Probe UDP traffic dropping by exploiting the DNS injection as a side channel
UDP-dropping IPs

Figure 3: An overview of the measurement pipeline of *IRBlock* with two scanning phases: the initial scans of the entire Iran's IP space to identify IP addresses subjected to DNS and HTTP filtering, and the subsequent scans that (1) distribute the testing of over 700M FQDN across the censored IP addresses identified in the initial scans to identify DNS and HTTP censored domains and (2) use the list of IP addresses subjected to DNS censorship to detect UDP traffic dropping.

4.1 Initial Scans of Iran's Entire IP Space

As shown in Figure 3, the first phase of our pipeline involves scanning the entire Iran's IP space to identify IP addresses subjected to DNS and HTTP censorship. We conduct these scans daily to ensure that our data is up-to-date and to capture any changes in the GFI's blocking behavior. This first step is crucial as we notice that the GFI's blocking behavior is not consistent across all IP addresses. Some IP addresses do not experience censorship at all.

For instance, sending a DNS query carrying a known censored domain (e.g., `twitter.com`) to `185.143.234.120`, which hosts the President of Iran's website (`president.ir`), does not trigger any DNS injections. Similarly, two IP addresses—`109.201.19.184` and `109.201.27.67`—belonging to Iran's Ministry of Foreign Affairs do not trigger any interference while neighboring IPs are subject to bidirectional censorship.

Remarkably, both IPs, as well as one of its neighbors that is subject to bidirectional censorship (`109.201.27.66`), have been implicated in cyber espionage campaigns by Playful Taurus (APT15), a Chinese state-sponsored threat group active since at least 2010. These addresses appear to host command-and-control infrastructure or spoofed security-related domains (e.g., `mfaantivirus[.]xyz`, `pfs1010[.]xyz`) used to deploy an updated variant of the Turian backdoor [18]. This malware enables persistent remote access and secure communication via obfuscated SSL channels and a randomized command structure. The sustained connections between Iranian government IPs and Playful Taurus infrastructure suggest that several state networks may have been compromised. Whether such IPs are deliberately excluded from GFI filtering due to operational reasons, were targeted because they are excluded from GFI filtering, or are simply anomalies remains unclear. Nonetheless, this underscores that censorship enforcement is not monolithic—even among sensitive government-operated infrastructure.

These examples highlight that censorship is selectively deployed across networks and subnets. Without a full scan of Iran's allocated IP space, such inconsistencies would be overlooked. Scanning the entire IP space is therefore essential not only to prevent false negatives (i.e., misclassifying a domain as uncensored due to probing uncensored endpoints), but also to reveal operational exceptions and potential allowlists within the GFI. We further discuss the ethical considerations surrounding these scans in §6.1.

For this initial scan, we send 19 DNS and HTTP probes to each IP address in Iran's IP space, which consists of over 11M allocated IPv4 addresses. We obtain this list of IP addresses from RIPE [11], and the list of domains is manually crafted to include a diverse set of domains, including benign domains, known censored domains, and a domain under our control that we know is not censored. Our probes are sent from two clusters of machines located in an educational network in North America where we have control over the network infrastructure which we confirm no filtering policies are in place.

An IP is identified as subjected to DNS censorship if we observe a DNS injection after sending a DNS query for a known censored domain. Similarly, an IP is identified as subjected to HTTP censorship if we observe a HTTP blockpage injection after sending a HTTP request containing a known censored domain. Our initial scans started in November 2024, *IRBlock* is still actively conducting these scans as of this writing. For this paper, we present the results of our initial scans conducted over a period of 2.5 months from November 2024 to January 15, 2025.

4.2 Distributed Testing of Censored Domains

In the second phase, we distribute the testing of over 700M fully-qualified domain names (FQDNs) across the censored IP addresses identified in the initial scans to identify DNS and HTTP censored domains. This second phase of our measurement pipeline started in January 2025. As one of our design goals is not to rely on responsive IPs for our measurements, using auxiliary information provided by Censys [33], we exclude IP addresses that are known to be responsive from subsequent tests. This ensures that we do not generate unnecessary traffic to responsive IPs and minimizes the impact of our probing activities on these IP addresses.

Inspired by the design choices of recent large-scale censorship measurement systems [43, 44], our list of test domains is curated by combining different sources of domains and to mitigate the effect of potential biases in the domains included in each list, which would otherwise introduce a bias in our results. These domain test lists are collected from various sources, including top-level domains (TLD) zone files [6], the Citizen Lab test lists (CLTL) [13], the Tranco list [60], and the Common Crawl project [2]. We use a new domain list generated every day to ensure that *IRBlock* probes up-to-date domains. In total, *IRBlock* probed over 700M FQDNs originating from over 500M apex domains.

4.3 UDP Traffic Dropping Detection

To detect UDP traffic dropping, we leverage the GFI's DNS injection as a side-channel. After we could determine the list of IP addresses that are consistently subjected to DNS censorship, we then send a UDP packet to each of these IP addresses, followed by a DNS query containing a known censored domain. These two packets share the same tuple of (source IP, source port, destination IP, destination port). If we observe a DNS injection, it indicates that the UDP traffic is not dropped. Conversely, if we do not observe a DNS injection, it indicates that the UDP traffic is dropped as shown in Figure 2. This technique allows us to detect UDP traffic dropping without requiring a responsive IP and provides a scalable way to measure UDP-based censorship mechanisms of the GFI.

Instead of using arbitrary payloads for the UDP packet, we build it to resemble a QUIC Initial Connection packet. This design choice is motivated by the fact that Iran has lower QUIC adoption compared to the rest of the world [58]. And as QUIC is on the rise, we opt to use this protocol to study UDP traffic dropping. We also use a unique source port for each UDP probe to avoid any potential noise from residual censorship. We started this probing phase in the third week of December 2024, after we could confirm a consistent set of IPs that are subjected to DNS censorship.

4.4 Scanning Strategy

Prior studies have shown that some nation-state censorship systems exhibit residual censorship, where subsequent packets are also tampered with, regardless of the content in the payload [43]. To mitigate the impact of residual censorship on our measurements, we adopt a unique scanning strategy that ensures that each DNS probe is sent from a unique source port, and each HTTP probe is sent from a unique source port and destination port combination. This strategy allows us to match our probes to responses and avoid any potential noise from residual censorship. Moreover, it also provides an easy and effective mechanism for quickly matching responses with our original probes. Note that we could apply this strategy because the GFI's HTTP filter has been activated on all TCP ports (§2.1).

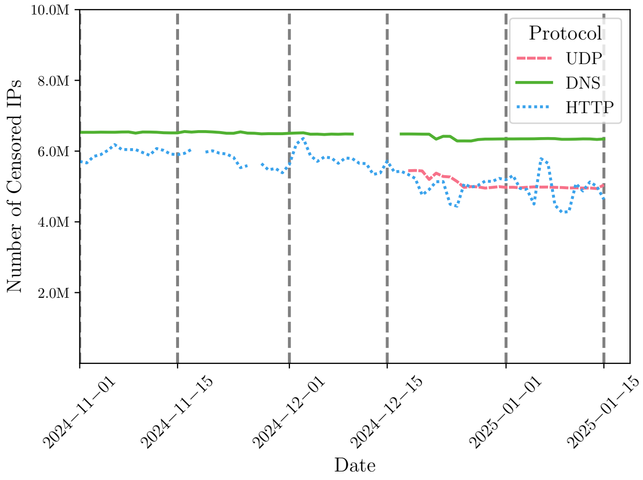

Figure 4: The number of IPs subject to DNS, HTTP, and UDP censorship over time. There are a few small gaps in the data due to the maintenance of our measurement machines. Nonetheless, the overall trend is visible and consistent.

This design choice also helps remedy the potential impact of Equal-Cost Multi-Path (ECMP) routing [27], where packets are distributed across multiple paths to the same destination. Each probe is also repeated at three different times of the day to account for potential packet loss and dynamic routing that may affect the consistency of our results.

5 Measurement Results

In this section, we present the results of our measurements and discuss in detail the differences between the observed censorship behaviors of the GFI's DNS, HTTP, and UDP filters. The data spans the period from November 2024 to January 15, 2025, during which we scanned all 11M Iranian IPs daily and tested over 700M fully qualified domain names (FQDNs). These longitudinal measurements uncover insights into the scope of censorship, temporal variations, and the operational strategies of the GFI. Specifically, we focus on three aspects: (1) which IPs are censored (§5.1), (2) what domains are blocked (§5.3), and (3) the blocking strategies of the GFI's injectors (§5.4), which we have found to share many similarities with strategies employed by the Great Firewall of China (GFW).

5.1 Censored IPs

Figure 4 shows the number of IP addresses that are subject to DNS, HTTP, and UDP censorship over time, where the censored IPs for each day are aggregated from all three probes on that day. To account for uncontrollable factors such as packet loss and noise, we only include IPs in our analysis for which we observe censorship at least three times.

DNS Trends. DNS censorship demonstrated remarkable stability throughout the measurement period, consistently affecting over 6.5M IPs daily until mid-December 2024. However, starting around the third week of December 2024, there was a gradual decline in censored IPs, reaching approximately 6.35M by January 2025. This decline aligns with reports suggesting that the Iranian government scaled back Internet restrictions during this period [38]. The stability of DNS blocking highlights its role as a foundational layer of the GFI's censorship infrastructure, with fewer variations compared to HTTP and UDP censorship.

HTTP Trends. HTTP censorship exhibited higher variability compared to DNS, with a slight but noticeable lag in its deployment. Although the overall number of censored IPs closely mirrored DNS trends, the fluctuations in HTTP censorship suggest possible constraints in the stateful HTTP filtering mechanisms. These could include resource limitations, rate-limiting policies, or operational inconsistencies within the GFI's infrastructure. In fact, from an operational perspective, HTTP censorship is more complex than DNS censorship, as it requires the inspection of packet payloads and the maintenance of stateful connections. Results from our analysis reveal that the overlap between DNS and HTTP-censored IPs exceeded 99%, further confirming that HTTP blocking is largely contingent on DNS censorship. Figure 4 also shows a slight decline in HTTP censorship around the same time as the drop in DNS censorship, indicating a coordinated effort to reduce censorship across both protocols. This trend further underscores the interdependence of DNS and HTTP censorship within the GFI's censorship infrastructure.

UDP Trends. UDP blocking, which we began measuring in mid-December 2024, affected a smaller subset of addresses compared to DNS and HTTP. The number of UDP-blocked IPs remained relatively stable at approximately 5M, with minor fluctuations. These IPs were exclusively a subset of those already censored via DNS since our method for detecting UDP traffic dropping relies on DNS censorship as a side-channel. We also notice a similar drop in UDP blocking around the same time as the decline in DNS and HTTP censorship. This is also correlated with the sharp rise in the proportion of HTTP/3 traffic [3] from AS58224, a large AS with heavy traffic interference (Table 1), which increased from 4.55% on December 16 to 10.01% on December 30. Even though we have no further quantitative evidence, this coincides with the announcement by the Iranian government scaling back Internet censorship [38] and the rise in HTTP/3 traffic could be attributed to the relaxation of censorship.

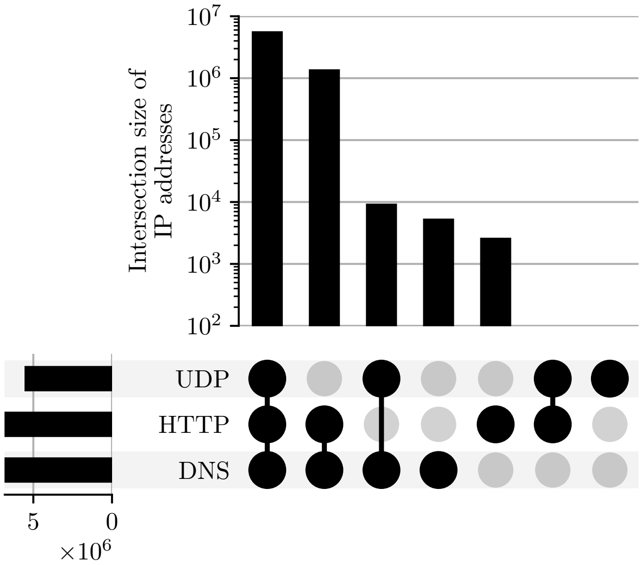

Figure 5: Correlation in IPs subject to DNS poisoning, HTTP blockpage injection, and UDP blocking.

Furthermore, when trying to correlate the censored IPs across all three protocols, as shown in Figure 5, we find that if an IP is subject to UDP blocking, it almost always also experiences HTTP and DNS blocking. This suggests that UDP blocking is used in conjunction with DNS and HTTP blocking. However, the opposite is not true, and about 1.4M IPs experience only DNS and HTTP blocking. This discrepancy could point to selective deployment of UDP filtering, potentially targeting specific networks or applications.

5.2 AS-level Blocking Behavior

AS Distribution of Censorship. As discussed earlier, one of the reasons why we design *IRBlock* to probe all Iranian IPs is to understand whether there are any patterns in the censorship behavior across different ASes. The case of networks being exempt from censorship discussed in §4.1 suggests a nuanced and possibly hierarchical structure of censorship deployment. Our analysis mapped censored IPs to their corresponding ASes, providing insights into how censorship is distributed and highlighting disparities in enforcement.

Using a combination of diverse datasets from RouteViews [12], DB-IP [4], ip-api [8], IPinfo [9], MaxMind [49], and RIPE NCC data [11] to determine the country and AS of an IP, and the country assignment of organizations combined

Table 1: DNS, HTTP, and UDP blocking behavior in all Iranian ASs with more than 100k assigned IPs

ASN	AS Name	AS Size	DNS			HTTP			UDP
			10.10.34.34	10.10.34.35	10.10.34.36	10.10.34.34	10.10.34.35	10.10.34.36	
58224	TCI	3.60M	89.66%	98.34% 65.80%	40.14%	0.19%	93.33% 65.77%	34.36%	87.61%
197207	MCCI	2.30M	0.76%	0.86% 0.72%	0.15%	0.00%	0.93% 0.72%	0.21%	0.15%
44244	IranCell	1.31M	0.48%	0.69% 0.69%	0.00%	0.00%	0.20% 0.20%	0.00%	0.69%
31549	RASANA	577k	97.33%	99.99% 23.38%	82.91%	0.06%	99.99% 23.38%	80.02%	64.05%
57218	RighTel	433k	0.64%	0.93% 0.87%	0.06%	0.00%	0.83% 0.77%	0.06%	0.78%
42337	RESPINA	262k	98.99%	99.28% 3.12%	98.49%	0.00%	99.09% 2.68%	98.18%	84.87%
39501	NGSAS	208k	32.12%	36.95% 36.95%	0.01%	0.17%	36.95% 36.95%	0.00%	36.95%
43754	ASIATECH	200k	86.66%	97.11% 77.01%	20.61%	0.20%	97.09% 77.00%	20.44%	26.01%
206065	FDI	153k	89.63%	98.71% 86.62%	35.20%	0.24%	98.70% 74.75%	27.76%	27.24%
50810	Mobinnet	130k	76.12%	87.88% 85.64%	2.80%	0.18%	87.88% 85.63%	2.76%	10.28%
24631	FANAPTELECOM	132k	97.88%	99.21% 13.32%	89.00%	0.04%	99.21% 11.56%	87.87%	76.31%
49100	IR-THR-PTE	112k	94.19%	99.97% 80.85%	59.54%	0.17%	99.95% 75.99%	58.99%	62.13%

with the inferred AS to organization data set by CAIDA [31] [2], we identified 816 ASes in Iran, of which 537 were active based on the criteria that an active AS has at least one prefix assigned to it. Of these 537 active ASes, 485 exhibited blocking behavior for at least 25% of their assigned IPs, underscoring the widespread deployment of censorship across the Iranian Internet infrastructure.

Table 1 provides an overview of DNS, HTTP, and UDP blocking behaviors for major ASes in Iran with more than 100K assigned IPs. The data reveals that some ASes, such as AS58224 and AS31549, apply nearly complete censorship across DNS, HTTP, and UDP protocols, whereas others, like AS197207 and AS44244, show negligible or no censorship. This disparity suggests that while the GFI operates as a centralized system, certain ASes may be granted exemptions or are subject to less stringent censorship measures.

Protocol-Specific Censorship. For DNS and HTTP, most censored ASes exhibit nearly identical levels of DNS and HTTP blocking. For instance, AS58224, the largest AS with over 3.6M assigned IPs, blocks more than 98% of IPs on both protocols. This trend aligns with the finding that DNS and HTTP censorship are almost always applied together. UDP blocking is less uniformly distributed. For example, AS43754 blocks 97% of IPs for DNS and HTTP but only 26% for UDP, suggesting that UDP censorship is selectively enforced. This selective enforcement may reflect prioritization of protocols based on perceived risk or operational complexity.

5.3 Censored Domains

After being able to confirm the consistency in the censored IPs across the different protocols, from December 2024 to January 2025, we tested over 700M Fully Qualified Domain Names (FQDNs) and 500M apex domains against DNS and HTTP censorship filters using *IRBlock*. Our findings reveal over 6M blocked FQDNs and 3.3M blocked apex domains. Among these, nearly 3M domains were blocked by DNS poisoning, and 1.6M by HTTP blockpage injection. To ensure accuracy of our analysis and eliminate transient one-off and short-term blocking cases, domains were only classified as censored if blocking occurred on at least three separate days.

Using public suffix data [10] and IANA's active TLD data [7], FQDNs are aggregated into apex domains to avoid inflating the number of blocked domains since technically a blocked apex domain (e.g., `hrw.org`) can have unlimited subdomains (e.g., `www.hrw.org`, `china.hrw.org`, `donate.hrw.org` etc.). Interestingly, 1.27M apex domains

[2] We have to use several different data sets because individual datasets are incomplete and inaccurate. Accurately identifying the AS and the country requires combining them.

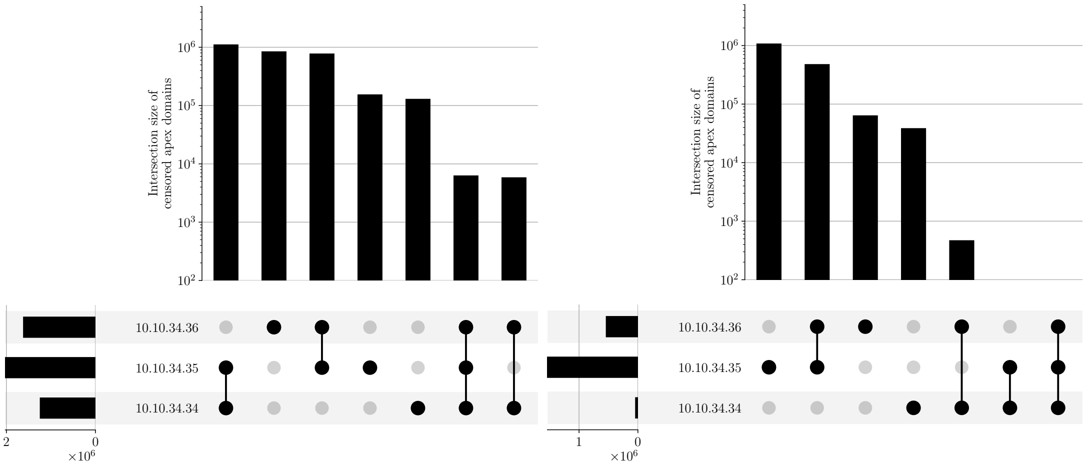

(a) DNS

(b) HTTP

Figure 6: Intersection of censored DNS and HTTP apex domains across GFI injectors (`10.10.34.34`, `10.10.34.35`, `10.10.34.36`), highlighting selective and overlapping censorship strategies.

were jointly censored by both DNS and HTTP filters, while significant discrepancies highlight distinct blocklists used by the two filtering methods. This underscores the operational independence of DNS and HTTP blocklists within the GFI.

Blanket Suffix Blocking. We observed aggressive blocking rules targeting entire top-level domains. Some blocking patterns were expected, such as bans on adult content domains (e.g., `.*\.porn$`). However, other patterns reveal unexpected blanket bans, including rules like `.*\.il$` (targeting all Israeli domains) and `.*\.com\.mx$` and `.*\.my\.id$`. These large TLD and suffix-level bans contribute to significant collateral damage, inadvertently blocking many legitimate websites. Of the 3.3M blocked apex domains, 1.7M were attributed to such suffix-level rules.

Censored Domain Categories. To better understand the nature of censored content, we categorized 87K apex domains ranked on the Tranco list [60] since this list provides a comprehensive overview of the most popular websites on the Internet. Focusing on the most popular domains allows us to identify the types of websites most impacted by censorship.

Using VirusTotal's domain classification service [14], we found that 37% of censored domains fell into the `adult content` category, followed by `entertainment` and `gambling`. The `Other` category contains both domains from less frequent content categories and those we could not classify (about 13.4k classified as `Unknown`). Blanket censorship rules discussed above were excluded from this analysis since it is imprecise to categorize entire TLDs (e.g., `.*il$`). Figure 7 summarizes the top domain categories, providing insight into the GFI's prioritization of censorship targets.

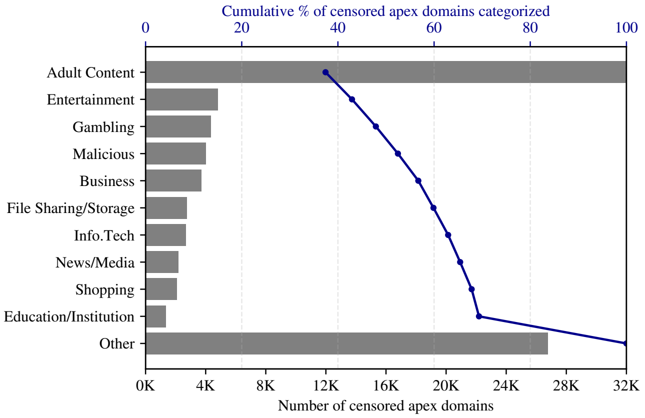

Figure 7: Top categories of the censored apex domains ranked on the Tranco list.

5.4 Strategic Similarities with China's GFW

Comparing the strategic similarities of the GFI and GFW, summarized in Table 2, we observe similarities in the general censorship behavior and blocking strategies. At the same time, the blocking behavior for the different protocols shows differences in the implementation.

Distinct Injectors and Blocking Patterns. As described in §2.1, we identified three unique injectors in the GFI infrastructure: `10.10.34.34`, `10.10.34.35`, and `10.10.34.36`. Each exhibits distinct blocking behaviors, similar to the GFW's triplet censor setup described in prior research [22, 44]. These injectors vary in their response mechanisms and target different sets of domains. While all three are involved in DNS and HTTP filtering, `10.10.34.35` plays a notable role in TTL

Table 2: Comparison of DNS, HTTP, HTTPS, and UDP blocking behaviors between the GFW and the GFI.

	The Great Firewall of China (GFW)	**The Great Firewall of Iran (GFI)**
Common Behavior	• Bidirectional censorship → enables remote measurement from outside of the country • Semi-stateful and TCP non-compliant → enables remote measurement without in-country volunteers / machines	
DNS	Fake IP injections [22, 42]: • Publicly routable addresses • Dynamic change (pool of ∼2K IPv4 addresses) • 3 different injectors • One DNS injector has a TTL reflection behavior [22]	Fake IP injections [23]: • Private IP addresses (10.10.34.x), but routable inside Iran • No dynamic rotation of injected IPs • 3 different injectors • One DNS injector has a TTL reflection behavior
HTTP	3x RST/ACK packets [22, 42]	Blockpage [23]
HTTPS	3x RST/ACK packets [22, 42]	1x RST/ACK packet [25, 62] (observed but not investigated at scale as part of this work)
UDP traffic dropping	Unknown	Traffic dropped based on the (source IP, source port, destination IP, destination port) tuple of the banned UDP flow

reflection-based injection, a behavior that was previously also reported in the GFW [22].

The presence of multiple injectors aligns with our observation that censorship behavior varies based on the destination IP, determined by the injector through which packets are routed. When categorizing censored Iranian IPs based on the injected IP, we observe a distinct partitioning, as illustrated in Figure 8. Notably, no IP exclusively receives injections from `10.10.34.34`, suggesting that this injector operates as a smaller, more selective component, likely triggered under specific conditions and separate from the two primary injectors. The main injectors, `10.10.34.35` and `10.10.34.36`, handle the majority of censorship, with all packets to censored IPs passing through at least one of these two. As depicted in Figure 6, the different injectors have substantially different blocklists, with each injector having a set of domains censored only by that injector. Some IPs even receive injections from both primary injectors, indicating that the censorship experienced by users can depend on packet routing or their geographic location. Table 1 provides a detailed breakdown of this partitioning, showing that, for Iranian ASs with more than 100K assigned IPs, the majority are affected predominantly by either `10.10.34.35` or `10.10.34.36`. While `10.10.34.34` impacts nearly all IPs during DNS queries, its involvement in HTTP probing is limited to a small subset of IPs.

In addition to the observed behavior of the three injectors, we identified a peculiar anomaly in the censorship behavior of the GFI. For DNS queries targeting google.com, the GFI interferes and injects a response IP; however, the injected IP surprisingly belongs to Google itself. We, thus, do not consider this case as DNS poisoning since the rationale behind this unusual behavior remains unclear.

TTL Reflection in Injected Responses. A hallmark of GFW's DNS censorship is the use of TTL reflection, where the TTL value of the probe packet is mirrored in the injected response [22]. During our experiments, the injector of `10.10.34.35` demonstrated identical behavior. For instance, when probing with a TTL of 23, the injected response reached us with a TTL of 1. Incrementing the probe TTL to 24 resulted in a response TTL of 2, and so on. This behavior necessitates doubling the TTL value of our probes to ensure the injected response traverses back through the network.

The purpose of this reflection remains unclear but serves as a potential strategy to obscure the censor's location. It complicates localization attempts using traditional TTL-limited probing methods, similar to the operational obfuscation tactics employed by the GFW.

Another strategic similarity between the GFI and the GFW is the partitioned domain blocklists where different injectors operate distinct blocklists that partially overlap, reflecting policy segmentation or infrastructure redundancy. These shared characteristics highlight the evolution of the GFI into a complex, multi-layer censorship system, comparable to the GFW. Understanding the parallels and differences enhances our ability to predict, analyze, and counteract similar censorship mechanisms in other nation-state firewalls.

6 Discussion

Next, we reflect on the ethical considerations and limitations of our study, as these aspects are crucial to conducting responsible and impactful research in a sensitive domain like censorship measurement. The ethical considerations involve evaluating the potential risks and benefits of our methodology, especially in relation to affected parties such as end users and on-the-ground volunteers. We also address the limitations inherent in our approach, including the constraints imposed by external measurements and the inability to fully explore certain censorship mechanisms due to the complexity of the GFI. By analyzing these aspects, we aim to provide a trans-

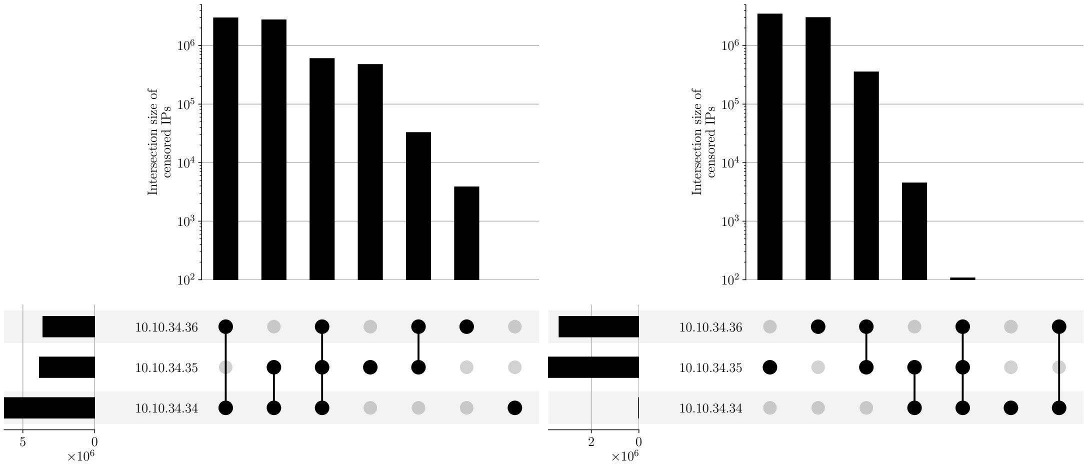

(a) DNS

(b) HTTP

Figure 8: Intersection of Iranian IPs subjected to DNS poisoning and HTTP blockpage injection across GFI injectors, showing partitioned and overlapping filtering behavior.

parent account of the decisions made during the development and execution of our study, as well as outline opportunities for future work to expand and refine our findings.

6.1 Ethics Considerations

Internet censorship research, particularly at the national scale, raises distinct ethical challenges given the political sensitivity of the topic and the risk of misattribution or unintended disruption [46, 66]. We assess our methodology through the lens of the ethical principles articulated in the Menlo Report [24] (*Respect for Persons, Beneficence, Justice, and Respect for Law and Public Interest*), guidance from USENIX Security regarding research ethics [16], and recent literature [47].

Tradeoffs and Motivations. Our primary goal is to measure the design and deployment of Iran's national censorship infrastructure while minimizing risk to individuals inside the country. To this end, *IRBlock* is designed to rely only on external probing techniques. We do not involve or intend to use volunteers, VPN infrastructure, or in-country public servers, as done in some prior measurement studies (e.g., OONI [37], ICLab [51], or Censored Planet [63]), thereby minimizing ethical exposure for local actors.

Instead, *IRBlock* uses bidirectional measurement techniques, relying solely on censorship infrastructure reacting to externally generated probes. This approach builds on and significantly extends prior work, such as Iris [59], Censored Planet [63], and TMC [54], which justify scanning public infrastructure for public-interest research while minimizing risk. Our design incorporates additional safeguards such as responsive-IP exclusion in the second phase of our measurements and opt-out mechanisms.

Impact of Our Measurements. We use a two-phase measurement process. The first phase involves lightweight probing of Iran's IPv4 address space to identify IPs subject to censorship. These probes—limited to 19 domains tested per IP—are minimal in bandwidth and designed to not initiate any actual application-layer sessions due to the unsolicited and stateless nature of our probing packets. These probes are sent to all IPv4 addresses, potentially including responsive ones. While this is not our intention, sending packets to responsive addresses introduces additional ethical complexity. However, omitting all responsive IPs, over 600k in Iran according to Censys data [33], would reduce the fidelity of our data and obscure important findings. For example, we would not have discovered stable exceptions like `president.ir` (mapped to `185.143.233.120`), or special uncensored IPs linked to state-sponsored malware operations [18]. As discussed in §4.1, these findings are crucial for understanding the granularity and policy exceptions of censorship deployment.

To mitigate any potential impact on responsive addresses, we have implemented several safeguards. First, one of the 19 domains tested is a domain under our control that we know is not censored and has the form `optout.our_research_domain.tld`. The presence of this domain in the payload would allow any entity observing it to reach out to us for an opt-out request. Second, our scanning machines have HTTP port 80 open to serve the opt-out page while their IPs are mapped to DNS PTR records, pointing to the opt-out domain as well. This is the best practice and widely adopted within Internet measurement research that involves network scanning [15, 34, 35].

Another potential ethical consideration is that our initial DNS scan of Iran's IPv4 address space may result in some queries for blocked domains being received by open DNS forwarders [52]. In such cases, the forwarder may relay the query to an upstream resolver, possibly generating an additional DNS lookup. Nonetheless, the forwarded request does not initiate any subsequent application-level communications with the censored domain (e.g., Web requests), and no end-user is directly involved. While we assess the likelihood of harm to be extremely low, we recognize that future studies may further reduce exposure by pre-filtering active forwarders using non-sensitive control domains. Future work can adopt such refinements to further reduce ethical risk while maintaining measurement fidelity. In retrospect, to eliminate even this small risk we could have excluded the about 11k IPs in Iran that respond to DNS queries on port 53 (Censys, May 2025 [33]). The impact of excluding those DNS-responsive IPs from our initial DNS scans would not have affected our discovery of special-exception IPs, since none of them answered on port 53.

For the second measurement phase, we only probe IPs that have been identified as censored in the first phase and are not alive based on up-to-date data provided by Censys [33]. Moreover, we have revamped our measurement tools to include active SYN checks before sending more probes, minimizing unnecessary contact with responsive IPs if Censys data happens to be outdated or missing some.

Finally, we note that if Iranian authorities were to investigate our probes, it would be evident that all measurement traffic originates from our machines located outside of Iran. At no point do we rely on infrastructure or relays within the country. This design decision was intentional and reflects our goal of minimizing risk to individuals inside the country. As of this writing, we have not received any complaints or opt-out requests related to our probing activities, suggesting that the traffic did not cause any disruptions or has been misinterpreted by network operators.

Testing Frequency Considerations. Given that the majority of censored infrastructure in Iran appears to not change drastically over the course of our study, we can reduce scan frequency for long-term monitoring from daily to once per week or spread it over a longer time frame. This reduction further minimizes exposure while maintaining the ability to detect policy shifts or infrastructure reconfiguration. However, it is important to note that `IRBlock` is designed to flexibly launch new scans any time granularity in response to emerging events, if needed.

Public Benefits. In accordance with the Menlo Report's principle of *Justice* [24], our research is designed to support a broad range of stakeholders—including researchers, journalists, policymakers, and circumvention tool developers. By identifying blocking infrastructure, policy granularity, and operational fingerprints of censorship in Iran, we aim to support digital rights advocacy and informed public debate.

We will publish our dataset and measurement scripts, enabling transparent validation of our study and facilitates future research. One may concern that releasing our work could also help the censor to improve their censorship infrastructure. However, we believe that the benefits of transparency and open science outweigh the risks. The Iranian government is already aware of the existence of censorship measurement and circumvention tools. By sharing our findings, we aim to empower individuals and organizations working to resist censorship and promote digital rights.

To that end, `IRBlock` is designed to balance the imperative to understand widespread censorship with the obligation to minimize risk. We strive to be transparent, proactive, and responsive to feedback from the community. We are committed to keeping `IRBlock` operational and to continuously improving our measurement techniques. We believe that the potential benefits of our work—in documenting and enabling resistance to repressive information controls—outweigh the minimal and carefully mitigated risks associated with our methodology.

6.2 Limitations

While `IRBlock` has made significant strides in understanding the GFI, certain limitations remain, primarily due to the complexity of the system and the constraints of external-only measurements.

Other Blocking Mechanisms. Our study focuses on DNS, HTTP, and UDP-based censorship, while omitting HTTPS filtering due to technical and ethical challenges. HTTPS blocking often involves silent packet drops, making it difficult to reliably detect at scale [25, 62]. Additionally, mechanisms like protocol whitelisting, previously observed in Iran [29], were not included in our analysis. Although these areas remain unexplored, our focus on the GFI's major protocols fills significant gaps in existing research and provides a strong foundation for future studies.

Lack of In-Country Vantage Points. One key property of `IRBlock` is that it does not require vantage points inside Iran for any of its measurements. The trade-off with this design is that we are unable to report on the packets sent by the GFI to the target IPs that we probe and therefore, we cannot detect any additional censorship mechanisms potentially only applied to traffic originating from inside Iran. Future research leveraging in-country vantage points could complement our findings by providing deeper insights into regional variations in blocking strategies. Together with other platforms like Censored Planet [63] and OONI [37], we hope to inspire further research in this direction, with each platform providing a unique perspective on the GFI, and when combined, offering a comprehensive view of the censorship landscape in Iran.

Potential Measurement Interference and Discrimination. A potential risk is that the measurement traffic by `IRBlock` is detected by the GFI and treated differently to manipulate our measurement results, as our probes stick out as abnormal Internet traffic. As far as we can tell from our measurements, we did not find any indication that the GFI is actively interfering with our measurements. While we already make an effort to randomize many dimensions of our probing, this is an active research area and IRBlock can easily be expanded to employ advanced techniques for disguising measurement traffic, such as splitting the traffic between more machines, strategic pauses and switching to non-measurement network traffic in randomized intervals [19]. One key advantage of our design without vantage points inside Iran is our ability to easily move our measurement infrastructure to a different network or a server. We will continue to actively monitor our measurements for signs of interference, and take appropriate countermeasures if necessary.

Another potential risk is that the GFI could change its behavior in response to our measurements, e.g., by becoming fully stateful. However, we believe that such changes would require a non-trivial effort and investment from the GFI, as the current design choices are widely regarded as a "feature" to make the GFI robust to packet loss and different packet routing [26, 27], which can also be observed for other censorship systems [43, 54]. Similarly, the blocking of our UDP blocking measurement would require a more expensive inspection of traffic and stateful behavior from the GFI.

Overall, despite these limitations, IRBlock represents a major advancement in the field of censorship measurement, offering the most comprehensive study of the GFI to date. Our efforts underscore the importance of ethical research practices in sensitive domains and highlight opportunities for future work to address existing gaps.

7 Related Work

Measurement of Iran's Censorship. Prior studies have explored Iran's censorship mechanisms, albeit with limited scope and coverage. Aryan et al. [23] conducted an early analysis using in-country machines, testing approximately 9,000 domains and finding 14 censored domains via DNS and 2,000 via HTTP and HTTPS. Similarly, comparing data collected over the same period as our study, OONI [37] have relied on volunteers for measurements, testing 2.7K domains, of which 1.3K have indications of censorship. Censored Planet [63] has also conducted measurements using public responsive servers, testing 1.9K, 2.1K, and 2.1K domains against the DNS, HTTP, and HTTPS filters, respectively, finding 407, 538, and 535 domains experiencing high 'unexpected' blocking rates of more than 80%.

While valuable for longitudinal studies across multiple countries, such platforms are constrained by their dependency on publicly accessible infrastructure, limiting their domain testing and network reach.

In contrast, IRBlock achieves unprecedented scale by probing the entire Iranian IP address space, testing over 500 million apex domains and revealing nearly 3M and 1.63M domains censored by the DNS and HTTP filters, respectively. The large-scale nature of IRBlock's dataset demonstrates the GFI's aggressive censorship strategies, which were previously underreported. Furthermore, while earlier studies have noted DNS and HTTP blocking primarily on protocol default ports, IRBlock uncovers censorship applied across all TCP ports by the HTTP filter, suggesting recent changes in the GFI's capabilities and operational strategies.

Comparative Analysis. The Great Firewall of China (GFW) remains the most studied Internet censorship system, with efforts like GFWeb [43] testing over a billion domains exploiting the GFW's bidirectional blocking and sophisticated filtering mechanisms. In comparison, IRBlock reveals striking similarities with the GFW, including TTL reflection (see §5.4) and the deployment of multiple injectors for censorship. However, unlike the on-path GFW, our study reveals significant evidence that GFI operates as an in-path system capable of dropping packets, requiring novel measurement methodologies. IRBlock's findings also suggest a more fragmented and selective approach in Iran, such as injector-specific censorship behaviors across DNS and HTTP protocols.

The Turkmenistan's Censorship study by Nourin et al. [54] (TMC) provides another reference point for comparisons. Conducted entirely from outside the country, TMC tested over 15M domains across 22.7K IPv4 addresses, identifying 122K censored domains. While TMC demonstrated innovative methods for remote measurements, IRBlock surpasses it in scale, probing more than 11M IPv4 addresses and testing 700M million domains. Additionally, IRBlock uncovers silent traffic dropping, a nuanced censorship behavior absent from prior studies in Turkmenistan.

Advancing Censorship Measurement. IRBlock addresses key limitations of previous studies by combining the methodological rigor of large-scale probes with ethical safeguards to minimize risks for local users. Our approach complements global platforms like Censored Planet and OONI while extending the scale and granularity of measurements. This is evident in the comparison with OONI, Censored Planet, and Aryan et al. [23], where IRBlock's can help with improving testing coverage of DNS and HTTP censorship, but other platforms like Censored Planet and OONI can provide more detailed insights into HTTPS censorship. Indeed, recent work by Nourin et al. [53] demonstrates that similar non-endpoint-based measurement methods can effectively measure both HTTP and HTTPS censorship across numerous countries and even in IPv6 networks that are difficult to measure with existing approaches that require active vantage points inside the country.

To that end, by providing a detailed analysis of Iran's multi-protocol censorship system, IRBlock bridges several critical gaps in our understanding of the GFI and lays the groundwork for improving existing censorship measurement platforms.

Study	Measurement Method	Duration	DNS	HTTP	HTTPS	Coverage
		MM/YY	(Censored domains/Tested domains)			
OONI [37, 56]	Volunteers (end users)	11/24–01/25	1.3K/2.7K			40
Censored Planet [17, 63]	Public responsive servers	11/24–01/25	407/1.9K	538/2.1K	535/2.1K	70
Aryan et al. [23]	In-country machines	04/13–05/13	14/9K	2K/9K	2K/9K	1
`IRBlock`	No active vantage points	11/24–01/25	2.99M/500M	1.63M/500M	Null	537

Table 3: Comparison with prior censorship measurement studies, showing how `IRBlock` can complement existing efforts by providing a more comprehensive view of Iran's Internet censorship. Note that we do not conduct large-scale HTTPS censorship measurements in this paper as it is challenging and ethically unsound to do so without active vantage points inside the country.

8 Conclusion

In this paper, we presented `IRBlock`, a large-scale, multi-protocol measurement study of the GFI. Using the injection behavior of the GFI as a side-channel, we present a novel UDP blocking measurement technique that enables censorship measurements without vantage points. By applying this new technique along with other techniques using the bidirectional behavior of the GFI, `IRBlock` was able to overcome the major measurement challenge faced by past efforts, enabling us to conduct consistent scans of the entire Iranian IPv4 address space and to test the blocking status of over 500M apex domains. Our study presents the most comprehensive view of Iran's Internet censorship to date, revealing over 6M FQDN and 3.3M apex domains affected by DNS and HTTP injections across more than 6.8M IPs during a period of 2.5 months, along with 5.4M IPs affected by UDP-based filtering. Our findings reveal that blocking rules are not applied homogeneously but that the GFI consists of three distinct injectors with different blocking rules.

By providing the first large-scale, longitudinal study of the GFI, `IRBlock` complements existing censorship measurement and reveals detailed insights into the behavior and evolution of the GFI that would not have been possible with existing methods relying on vantage points inside the country. In measuring Iran, a country that is notoriously difficult to study, `IRBlock` presents a significant step forward in Internet censorship measurement, as our architecture and measurement methods have the potential to be adapted for other censorship systems.

We are committed to continue to provide regular updates on the GFI's censorship behavior and to hope that the data collected by `IRBlock` fosters further research on nation-state censorship and supports other Internet freedom initiatives. More recent data collected after this publication is available at `https://IRBlock.org`. Our efforts not only advance censorship measurement methodologies but also inform policymakers and activists working to promote digital freedom in Iran and beyond.

9 Open Science

As researchers in censorship measurement working towards an open Internet, we are committed to making our research artifacts publicly available. To stimulate reproducibility and replicability of our findings, we make all data collected by `IRBlock` publicly available at `https://doi.org/10.5281/zenodo.15572895`. This includes the dataset of censored domains, the IP addresses that experienced DNS poisoning, HTTP blockpage injection, and UDP-based traffic disruption. In addition, we also share the measurement scripts that can be used to trigger the censorship mechanisms of the Great Firewall of Iran (GFI).

Moreover, while the data presented in this paper is from November 2024 to January 15th, 2025, we will continue to collect data and make it available to the research community. We will provide regular updates on the GFI's censorship practices, including the identification of censored IPs and domains, the discovery of new blocking mechanisms, and the analysis of the GFI's blocking rules and behavior. We hope that this data will be useful for researchers, policymakers, and activists working to promote digital freedom in Iran and beyond.

Acknowledgments

We would like to thank the anonymous reviewers and shepherd for their thorough feedback on this paper. We also thank others who preferred to remain anonymous for helpful discussions and suggestions.

This work was supported by a fellowship of the German Academic Exchange Service (DAAD), the Internet Society Pulse Fellowship, and funding from the Natural Sciences and Engineering Research Council of Canada (NSERC) under awards RGPIN-2025-04929 and DGECR-2025-00173. This research was supported in part through computational resources and services provided by Advanced Research Computing at the University of British Columbia (UBC ARC). We also acknowledge the invaluable support provided by the UBC Department of Computer Science technical staff and the UBC Cybersecurity team. We also gratefully acknowledge travel support from the Institute for Computing, Information and Cognitive Systems (ICICS) at UBC, which provided funding through its conference and member support programs. The opinions in this paper are those of the authors and do not necessarily reflect the views of the sponsors.

References

[1] A summary on Iran's current internet situation. https://github.com/net4people/bbs/issues/182.

[2] Common Crawl Project. https://commoncrawl.org.

[3] Data explorer | cloudflare radar - http versions time series for as58224. https://radar.cloudflare.com/explorer?dataSet=http&dt=12w&loc=58224&groupBy=http_version.

[4] DB-IP: IP Geolocation API & Free Address Database. https://db-ip.com.

[5] Dynamic UDP port blocking. https://github.com/net4people/bbs/issues/181.

[6] ICANN Centralized Zone Data Service. https://czds.icann.org.

[7] Internet Assigned Numbers Authority: TLDs alpha per domain. https://data.iana.org/TLD/.

[8] IP Gelocation API. https://ip-api.com.

[9] IPinfo: The Trusted Source For IP Address Data. https://ipinfo.io.

[10] Public Suffix List. https://publicsuffix.org/.

[11] RIPE NCC: RIPEstat - Providing open data and insights for Internet resources. https://stat.ripe.net/about.

[12] Route-Views Data. https://www.routeviews.org/data.html.

[13] The Citizen Lab Test Lists. https://github.com/citizenlab/test-lists.

[14] VirusTotal: URL Scanning Service. https://www.virustotal.com/gui/home/url.

[15] Zmap - Scanning Best Practices — github.com. https://github.com/zmap/zmap/wiki/Scanning-Best-Practices.

[16] USENIX Security '25 Ethics Guidelines, June 2024.

[17] Censored Planet Dashboard, Accessed: 2025-01-22. https://dashboard.censoredplanet.org/access.html.

[18] Paolo Alto Unit 42. Chinese Playful Taurus Activity in Iran — unit42.paloaltonetworks.com. https://unit42.paloaltonetworks.com/playful-taurus/, 2023.

[19] Abderrahmen Amich, Birhanu Eshete, Vinod Yegneswaran, and Nguyen Phong Hoang. Deresistor: Toward detection-resistant probing for evasion of Internet censorship. In *Proceedings of the 32nd USENIX Security Symposium (USENIX Security '23)*, pages 2617–2633, 2023.

[20] Collin Anderson. The hidden Internet of Iran: Private address allocations on a national network. *arXiv preprint arXiv:1209.6398*, 2012.

[21] Collin Anderson. Dimming the internet: Detecting throttling as a mechanism of censorship in Iran. *arXiv preprint arXiv:1306.4361*, 2013.

[22] Anonymous, Arian Akhavan Niaki, Nguyen Phong Hoang, Phillipa Gill, and Amir Houmansadr. Triplet censors: Demystifying Great Firewall's DNS censorship behavior. In *Proceedings of the 10th USENIX Workshop on Free and Open Communications on the Internet (FOCI '20)*, 2020.

[23] Simurgh Aryan, Homa Aryan, and J. Alex Halderman. Internet censorship in Iran: A first look. In *Proceedings of the 3rd USENIX Workshop on Free and Open Communications on the Internet (FOCI '13)*, 2013.

[24] Michael Bailey, David Dittrich, Erin E. Kenneally, and Douglas Maughan. The Menlo report. *IEEE Security & Privacy*, 10:71–75, 2012.

[25] Simone Basso. Measuring sni based blocking in iran, 2020-04-28. https://ooni.org/post/2020-iran-sni-blocking/.

[26] Abhishek Bhaskar and Paul Pearce. Many roads lead to Rome: How packet headers influence DNS censorship measurement. In *Proceedings of the 31st USENIX Security Symposium (USENIX Security '22)*, pages 449–464, 2022.

[27] Abhishek Bhaskar and Paul Pearce. Understanding routing-induced censorship changes globally. In *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 437–451, 2024.

[28] Kevin Bock, Abdulrahman Alaraj, Yair Fax, Kyle Hurley, Eric Wustrow, and Dave Levin. Weaponizing middleboxes for TCP reflected amplification. In *Proceedings of the 30th USENIX Security Symposium (USENIX Security '21)*, pages 3345–3361, 2021.

[29] Kevin Bock, Yair Fax, Kyle Reese, Jasraj Singh, and Dave Levin. Detecting and evading censorship-in-depth: A case study of iran's protocol whitelister. In *Proceedings of the 10th USENIX Workshop on Free and Open Communications on the Internet (FOCI '20)*, 2020.

[30] Kevin Bock, George Hughey, Louis-Henri Merino, Tania Arya, Daniel Liscinsky, Regina Pogosian, and Dave Levin. Come as you are: Helping unmodified clients bypass censorship with server-side evasion. In *Proceedings of the Annual conference of the ACM Special Interest Group on Data Communication on the applications, technologies, architectures, and protocols for computer communication (SIGCOMM '20)*, pages 586–598, 2020.

[31] CAIDA. Inferred as to organization mapping dataset, Accessed: 2025-01-15. `https://www.caida.org/catalog/datasets/as-organizations/`.

[32] Jakub Dalek, Adam Senft, Masashi Crete-Nishihata, Ronald J. Deibert, and Greg Wiseman. Behind blue coat: An update from Burma. 2011.

[33] Zakir Durumeric, David Adrian, Ariana Mirian, Michael Bailey, and J. Alex Halderman. A search engine backed by Internet-wide scanning. In *Proceedings of the 2015 ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 542–553, 2015.

[34] Zakir Durumeric, Michael Bailey, and J. Alex Halderman. An Internet-wide view of Internet-wide scanning. In *Proceedings of the 23rd USENIX Security Symposium (USENIX Security '14)*, pages 65–78, 2014.

[35] Zakir Durumeric, Eric Wustrow, and J. Alex Halderman. ZMap: Fast Internet-wide scanning and its security applications. In *Proceedings of the 22nd USENIX Security Symposium (USENIX Security '13)*, pages 605–620, 2013.

[36] Kathrin Elmenhorst, Bertram Schütz, Nils Aschenbruck, and Simone Basso. Web censorship measurements of HTTP/3 over QUIC. In *Proceedings of the 2011 ACM Internet Measurement Conference (IMC)*, pages 276–282, 2021.

[37] Arturo Filastò and Jacob Appelbaum. OONI: Open observatory of network interference. In *Proceedings of the the 2nd USENIX Workshop on Free and Open Communications on the Internet (FOCI '12)*, 2012.

[38] Christina Fincher. Iran lifts ban on WhatsApp and Google Play, state media says. Reuters. `https://www.reuters.com/technology/iran-lifts-ban-whatsapp-google-play-state-media-says-2024-12-24/`.

[39] Freedom House. Freedom on The Net 2024 - Iran, 2024. `https://freedomhouse.org/country/iran/freedom-net/2024`.

[40] Geremie R. Barme And Sang Ye. The Great Firewall of China, 1997-06-01. `https://www.wired.com/1997/06/china-3/`.

[41] Michael Harrity, Kevin Bock, Frederick Sell, and Dave Levin. GET /out: Automated discovery of application-layer censorship evasion strategies. In *Proceedings of the 31st USENIX Security Symposium (USENIX Security '22)*, pages 465–483, 2022.

[42] Nguyen Phong Hoang. GFWatch Dashboard, 2020. `https://gfwatch.org`.

[43] Nguyen Phong Hoang, Jakub Dalek, Masashi Crete-Nishihata, Nicolas Christin, Vinod Yegneswaran, Michalis Polychronakis, and Nick Feamster. GFWeb: Measuring the Great Firewall's Web censorship at scale. In *Proceedings of the 33rd USENIX Security Symposium (USENIX Security '24)*, pages 2617–2633, 2024.

[44] Nguyen Phong Hoang, Arian Akhavan Niaki, Jakub Dalek, Jeffrey Knockel, Pellaeon Lin, Bill Marczak, Masashi Crete-Nishihata, Phillipa Gill, and Michalis Polychronakis. How great is the Great Firewall? Measuring China's DNS censorship. In *Proceedings of the 30th USENIX Security Symposium (USENIX Security '21)*, pages 3381–3398, 2021.

[45] Jana Iyengar and Martin Thomson. RFC 9000: QUIC: A UDP-based multiplexed and secure transport. *Omtermet Emgomeeromg Task Force*, 2021.

[46] Ben Jones, Roya Ensafi, Nick Feamster, Vern Paxson, and Nicholas C. Weaver. Ethical concerns for censorship measurement. pages 17–19, 2015.

[47] Tadayoshi Kohno, Yasemin Gülsüm Acar, and Wulf Loh. Ethical frameworks and computer security trolley problems: Foundations for conversations. In *Proceedings of the 32nd USENIX Security Symposium (USENIX Security '23)*, pages 5145–5162, 2023.

[48] Ed. M. Bishop. HTTP/3. RFC 9114, IETF, June 2022.

[49] Maxmind. Industry leading IP Geolocation and Online Fraud Detection. `https://www.maxmind.com/en/home`.

[50] P. Mockapetris. Domain names - concepts and facilities. RFC 1034, IETF, November 1987.

[51] Arian Akhavan Niaki, Shinyoung Cho, Zachary Weinberg, Nguyen Phong Hoang, Abbas Razaghpanah, Nicolas Christin, and Phillipa Gill. ICLab: A global, longitudinal Internet censorship measurement platform. In *Proceedings of the 2020 IEEE Symposium on Security & Privacy*, pages 135–151, 2020.

[52] Arian Akhavan Niaki, William R. Marczak, Sahand Farhoodi, A. McGregor, Phillipa Gill, and Nicholas C. Weaver. Cache Me Outside: A New Look at DNS Cache

Probing. In *Proceedings of the 2021 Passive and Active Measurement Conference (PAM)*, 2021.

[53] Sadia Nourin, Erik Rye, Kevin Bock, Nguyen Phong Hoang, and Dave Levin. Is Nobody There? Good! Globally Measuring Connection Tampering without Responsive Endhosts. In *Proceedings of the 2025 IEEE Symposium on Security and Privacy (IEEE S&P '25)*, pages 1400–1418, 2025.

[54] Sadia Nourin, Van Tran, Xi Jiang, Kevin Bock, Nick Feamster, Nguyen Phong Hoang, and Dave Levin. Measuring and evading Turkmenistan's Internet censorship. In *Proceedings of the 2023 ACM Web Conference (WWW)*, pages 1969–1979, 2023.

[55] Office of Foreign Assets Control (OFAC). An overview of O.F.A.C. Regulations involving Sanctions against Iran.

[56] OONI. OONI Explorer: Uncover evidence of Internet censorship worldwide, Accessed: 2025-01-22. `https://explorer.ooni.org/`.

[57] Ramakrishna Padmanabhan, Arturo Filastò, Marianna Xynou, Ram Sundara Raman, Kennedy Middleton, Mingwei Zhang, Doug Madory, Molly Roberts, and Alberto Dainotti. A multi-perspective view of Internet censorship in Myanmar. pages 27–36, 2021.

[58] Imad Payande. Internet in Iran; censorship, sanctions, and the challenges facing users. *Censorship, Sanctions, and the Challenges Facing Users (June 01, 2024)*, 2024.

[59] Paul Pearce, Ben Jones, Frank Li, Roya Ensafi, Nick Feamster, Nick Weaver, and Vern Paxson. Global measurement of DNS manipulation. In *Proceedings of the 26th USENIX Security Symposium (USENIX Security '17)*, pages 307–323, 2017.

[60] Victor Le Pochat, Tom van Goethem, Samaneh Tajalizadehkhoob, Maciej Korczyński, and Wouter Joosen. Tranco: A research-oriented top sites ranking hardened against manipulation. In *Proceedings of the 2019 Network and Distributed System Security Symposium (NDSS)*, 2019.

[61] Ram Sundara Raman, Leonid Evdokimov, Eric Wustrow, J. Alex Halderman, and Roya Ensafi. Investigating large scale HTTPS interception in Kazakhstan. In *Proceedings of the 2020 ACM Internet Measurement Conference (IMC)*, pages 125–132, 2020.

[62] Ram Sundara Raman, Louis-Henri Merino, Kevin Bock, Marwan M. Fayed, Dave Levin, Nick Sullivan, and Luke Valenta. Global, passive detection of connection tampering. *Proceedings of the Annual conference of the ACM Special Interest Group on Data Communication on the applications, technologies, architectures, and protocols for computer communication (SIGCOMM '23)*, pages 622–636, 2023.

[63] Ram Sundara Raman, Prerana Shenoy, Katharina Kohls, and Roya Ensafi. Censored Planet: An Internet-wide, longitudinal censorship observatory. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 49–66, 2020.

[64] Reethika Ramesh, Ram Sundara Raman, Matthew Bernhard, Victor Ongkowijaya, Leonid Evdokimov, Anne Edmundson, Steven Sprecher, Muhammad Ikram, and Roya Ensafi. Decentralized control: A case study of Russia. In *Proceedings of the 2020 Annual Network and Distributed System Security Symposium (NDSS)*. The Internet Society, 2020.

[65] Will Scott, Thomas Anderson, Tadayoshi Kohno, and Arvind Krishnamurthy. Satellite: Joint analysis of CDNs and network-level interference. In *Proceedings of the 2016 USENIX Annual Technical Conference (ATC)*, pages 195–208, 2016.

[66] Simone Basso and Maria Xynou and Arturo Filasto. China is blocking OONI, 2023-07-28. `https://ooni.org/post/2023-china-blocks-ooni/`.

[67] Ram Sundara Raman, Adrian Stoll, Jakub Dalek, Reethika Ramesh, Will Scott, and Roya Ensafi. Measuring the deployment of network censorship filters at global scale. In *Proceedings of the 2020 Network and Distributed System Security Symposium (NDSS)*, 2020.

[68] Benjamin VanderSloot, Allison McDonald, Will Scott, J. Alex Halderman, and Roya Ensafi. Quack: Scalable remote measurement of application-layer censorship. In *Proceedings of the 27th USENIX Security Symposium (USENIX Security '18)*, pages 187–202, 2018.

[69] Diwen Xue, Benjamin Mixon-Baca, ValdikSS, Anna Ablove, Beau Kujath, Jedidiah R. Crandall, and Roya Ensafi. TSPU: Russia's decentralized censorship system. In *Proceedings of the 2022 ACM Internet Measurement Conference (IMC)*, pages 179–194, 2022.