Computer Science Review 62 (2026) 101002

Contents lists available at ScienceDirect

Computer Science Review

journal homepage: www.elsevier.com/locate/cosrev

A review of internet censorship: Modern measurement and circumvention techniques

Thomas Grübl[a,*], Francisco Enguix[b], Burkhard Stiller[a]

[a] Communication Systems Group CSG, Department of Informatics IfI, University of Zürich UZH, Binzmühlestrasse 14, Zürich, CH-8050, Switzerland
[b] Valencian Research Institute for Artificial Intelligence (VRAIN), Universitat Politècnica de València (UPV), Camino de Vera S/N, Valencia, 46022, Spain

HIGHLIGHTS

- Comprehensive analysis of 146 Internet censorship measurement and circumvention studies from a technological perspective.
- Classification of censorship measurement studies by network protocols, countries, duration, methods, and results.
- Classification of censorship circumvention studies by protocols, throughput, and application-agnosticism.
- Taxonomy of circumvention methods: routing-based and obfuscation-based tools.

ARTICLE INFO

Keywords:
Internet censorship measurement
Internet censorship circumvention
Network protocols
Security
Privacy

ABSTRACT

As Internet censorship continues to be deployed across a number of nation-states, understanding its scope and underlying mechanisms is more important than ever. Consequently, research on censorship measurement and circumvention has attracted growing academic interest, particularly in recent years. This article provides an overview of the current state of the art in the field of Internet censorship measurement and circumvention research. First, a brief overview of the fundamentals is provided, followed by an in-depth analysis of 146 contemporary Internet censorship measurement and circumvention studies, predominantly those published within the last ten years, by applying a semi-systematic literature review methodology. Subsequently, the review briefly summarizes the ethical considerations in the field, it visualizes the geographical focus of censorship measurement studies, and it provides an overview of Internet protocols used to measure censorship. In addition, it presents a taxonomy of censorship circumvention tools, analyzes their key characteristics, and examines the prevalence of the underlying network protocols used in circumvention tools. The findings suggest that, while there are numerous solutions for circumventing censorship, many are niche or theoretical, and their practicality remains unknown. Although there is an observable trend toward large-scale longitudinal censorship measurement studies, the real-world effectiveness of (academic) censorship circumvention methods is rarely evaluated. Since both censorship measurement and circumvention research go hand-in-hand, there is an increasing number of measurement studies which directly translate their findings into practical circumvention strategies.

1. Introduction

With the rapid increase in digital communication over the last decades, efforts to control and restrict the free flow of information have also been on the rise. The growing deployment of censorship techniques [1–3] as well as increasing global attention to digital rights and freedoms [4–6] have fundamentally shaped the field over the past decade.

Alongside censorship efforts, a variety of Internet censorship circumvention tools have emerged, which allow users to regain access to information by circumventing censorship ranging from localized and temporal restrictions to large-scale and long-term state-level interventions. Some tools are more effective and user-friendly than others, serving different levels of technical expertise. What was once the domain of tech-savvy users, human rights activists, and journalists in the early

* Corresponding author.
Email address: gruebl@ifi.uzh.ch (T. Grübl).

https://doi.org/10.1016/j.cosrev.2026.101002
Received 14 July 2025; Received in revised form 26 December 2025; Accepted 24 May 2026
Available online 3 June 2026

days of censorship circumvention, has now become accessible to a broader audience, in part due to the commercialization and widespread adoption of Virtual Private Network (VPN) services [7].

In addition to VPNs, encrypted overlay networks for anonymous communication, such as The Onion Routing (Tor) [8], The Invisible Internet Project (I2P) [9], or Hyphanet [10], play an important role in circumventing censorship. These networks route traffic over multiple distributed nodes and provide strong privacy and anonymity guarantees for their users.

According to the 2025 Freedom House Report [11], media freedom nowadays faces widespread challenges in over 120 countries and territories, which include arrests, violence, legal harassment, and censorship. Furthermore, the report identified a general decline in global freedom for the 19th consecutive year in 2024. This trend is also observable in the digital space, where Internet freedom is increasingly under threat. A growing number of governments have implemented different forms of Internet censorship in recent years [12]. In some regions, entire segments of the population have experienced partial or complete shutdowns of Internet services, particularly during periods of political unrest or elections [13–17]. Such government-imposed censorship is commonly implemented on an Internet Service Provider (ISP) level using firewalls, content inspection software, and the like. These tools allow authorities to block access to specific websites or online platforms. They can also filter based on network protocols, source or destination Internet Protocol (IP) addresses and ports, and Domain Name System (DNS) requests or the Transport Layer Security (TLS) Server Name Indication (SNI) field. In more extreme cases, encrypted traffic may be blocked altogether based on protocol fingerprints and statistical features of encrypted traffic; one such example was temporarily observed in China in November 2021 [18]. The goal of such restrictions is often to suppress dissent, control information flow, or prevent mobilization.

Since the rise of large-scale deployments of censorship technologies, the body of research around censorship measurement and circumvention has been growing consistently. This review paper consolidates Internet censorship measurement and circumvention research. It analyzes and contextualizes different forms of measurement and circumvention techniques. Both measurement and circumvention tools utilize all commonly-used application-layer protocols, such as HTTPS, DNS, SMTP, and SSH. Examples of measurement techniques include network layer, transport layer, and application layer probing, or various fuzzing strategies that modify the values of protocol-specific fields. Whereas some measurement studies purely assess whether censors are blocking certain protocols and services (*e.g.,* [19–22]), others also aim to identify the physical location of the censoring devices (*e.g.,* [23–26]), uncover the exact blocking behavior (*e.g.,* [23,27–30]), or verify if the blocking behavior remains consistent over longer periods of time (*e.g.,* [13,14,18,31–34]). Circumvention techniques comprise proxy-based systems, publisher-centric systems, protocol behavior imitation methods, covert tunneling, refraction networking, Alibi routing, as well as pure steganographic approaches.

For example, publisher-centric circumvention leverages Content Delivery Networks (CDNs) that host content across numerous global servers, making it difficult for censors to block specific sites without disrupting essential web services. Similarly, refraction networking (formerly decoy routing) requires the help of Internet Service Providers (ISPs) to intercept and reroute user traffic intended for blocked sites through covert pathways that circumvent censorship. Unlike the two aforementioned methods which rely on support from third parties, most other circumvention approaches can be implemented by end users directly or with the assistance of volunteer networks. Additionally, in recent years, some authors [2,35,36] have developed large-scale platforms, such as OONI [37], Censored Planet [38], and ICLab [39], that continuously perform censorship measurements around the globe and openly publish their datasets. These platforms have contributed significantly to advancing our understanding of Internet censorship practices and trends worldwide. They allow researchers to analyze censorship events in near real-time, study the evolution of blocking techniques over time, and compare censorship patterns across regions.

2. Contributions

This study focuses explicitly on the technical aspects of Internet censorship. In particular, how it is measured, how it evolves over time, and how it can be circumvented through technical means. This paper does not aim to make claims about where conventional (non-digital) censorship is implemented, how widespread or impactful it is, or to evaluate its political, legal, or moral implications. Although there are interconnected topics, such as surveillance and content moderation, they are not the primary focus of this paper unless they contribute directly to technical developments in censorship measurement or circumvention. Similarly, while “analog” or non-technical forms of censorship, like censorship in educational publishing [40] or instances of self-censorship [41], offer important insights, they fall outside the scope of this study. The specific inclusion and exclusion criteria for the literature selection process are outlined in detail in Section 4.2. In short, only peer-reviewed articles that focus on the technical aspects of Internet censorship were included. These include studies that measure censorship in real-world or simulated environments, propose new or improved circumvention tools, or apply security and privacy techniques in the context of censorship. Articles were excluded if they focused mainly on political, legal, or social issues without a strong technical component.

Fundamentally, research on Internet censorship can be grouped into offensive and defensive methods. Offensive methods are censorship techniques deployed by censors with the intent to control and restrict users from accessing certain resources. Defensive methods comprise active (*i.e.,* censorship circumvention) and passive (*i.e.,* non-interfering censorship measurements) techniques. This paper reviews the defensive aspects of Internet censorship. Many studies [23,25,42–46] propose new circumvention strategies based on the insights gained through measurement studies. The choice of the right measurement strategies can directly lead to new evasion strategies. Since measurement and circumvention techniques go hand in hand, this review paper is the first to provide a comprehensive synthesis of the two aspects.
This paper makes the following contributions:

- ✓ A classification and synthesis of Internet censorship measurement studies, encompassing the considered network protocols, countries, study durations, and a summary of the methods and results.
- ✓ A classification and synthesis of Internet censorship circumvention studies, focusing on the used network protocols, the throughput, evaluation strategies, application-agnosticism, and tool-specific features.
- ✓ A taxonomy of Internet censorship circumvention methods, identifying and reviewing the two primary categories: routing-based and obfuscation-based techniques.
- ✓ An analysis of the geographic focus of censorship measurement studies from the past decade, as well as a visualization of the key technical attributes of circumvention methods, including support for bootstrapping, operational readiness, and application agnosticism.

2.1. Paper outline

This paper begins with an overview of related work in Section 3, focusing on existing studies of Internet censorship from a technical

standpoint, as well as an analysis of the evolution of literature in the field. Section 4 presents the methodology, including the search strategy and criteria for the inclusion and exclusion of sources. Section 5 provides the necessary background by defining relevant terminology and offers a short introduction to censorship measurements and circumvention techniques. The core analysis begins in Section 6, which presents a thorough review of censorship measurement literature, including a discussion of ethical considerations in the field. Section 7 continues with an in-depth analysis of circumvention methods and provides a taxonomy to classify different circumvention approaches. Section 8 briefly discusses the interdependency of Internet censorship measurements and circumvention research and Section 9 makes remarks on the reproducibility of censorship measurement studies. Finally, Section 10 concludes the paper with a summary of key findings and a discussion of broader implications.

3. Related work

Since the field has rapidly grown in recent years (see Fig. 1), there are no existing survey papers that capture all the recent developments in Internet censorship measurement and circumvention research. Related survey papers are presented in Table 1. Some are narrowly focused, such as [47], a multi-perspective exploratory survey on the needs of circumvention tool providers, or [48], which analyzes the gap between practice and research. They argue that while forward-looking research on sophisticated future censorship is valuable, there is a critical need to focus more on practical, real-world censorship threats that are currently exploitable.

The authors of [49] provide an analysis of the design principles of six Future Internet Architectures (FIA), including but not limited to Named Data Networking (NDN), MobilityFirst (MF), and SCION (Scalability, Control, and Isolation On Next-Generation Networks). Such FIAs fundamentally differ from the "'traditional" Internet architecture in that they implement custom topological concepts, control plane and data plane functionalities.

A cross-sectional study of 70 countries is presented in [50], which includes an overview of different Internet censorship methods, the censorship trends over the past two decades, and a framework for ongoing censorship monitoring and analysis.

The authors of [51] specifically survey covert communication techniques, focusing on the underlying principles, methodologies, and practical applications. The review includes an in-depth analysis of radio frequency (RF)-based methods, traditional steganographic approaches applied to digital media (*e.g.,* images, videos, audio files, and text) as well as unconventional covert communication techniques, such as blockchain-based covert channels.

The authors of [53] introduce a framework to evaluate Censorship Resistance Systems (CRS) based on security, privacy, performance, and deployability. They provide an evaluation of various circumvention tools and their respective bootstrapping and operational phases. They also identify and describe key challenges, including the need to constantly reassess censor capabilities in changing sociopolitical contexts, as well

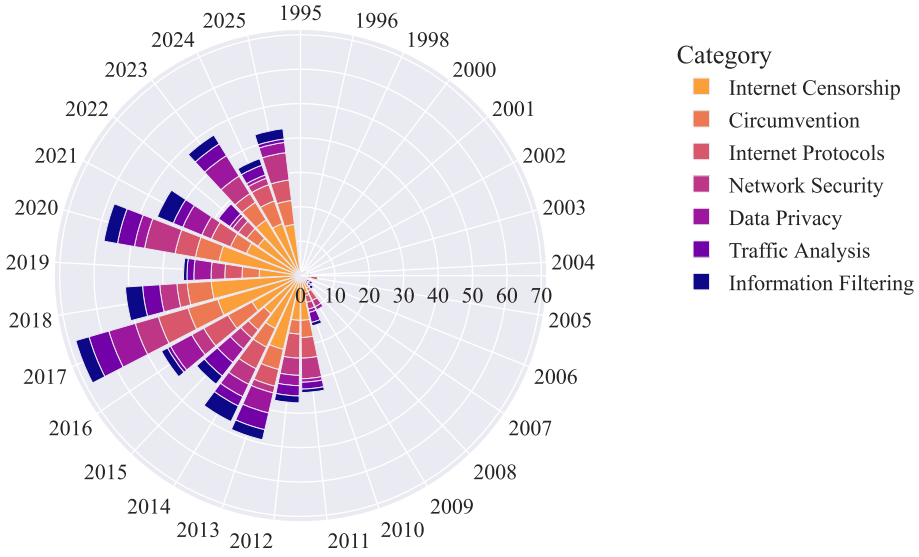

Fig. 1. Thematic Breakdown of research related to Internet censorship from 1995 to 2025, based on Scopus data. All categories include only papers that directly pertain to Internet censorship.

Table 1
Comparison of surveys on Internet censorship research.

Work	Year	Fundamentals	Measurements	Circumvention	Focus Area
[52]	2025	✔	✔	✘	Measurement methodologies & challenges.
[49]	2025	✔	✘	✘	Potential effectiveness of censorship in FIAs.
[47]	2024	✘	✘	✔	Challenges posed by circumvention tools.
[50]	2023	✘	✔	✘	Censorship trends in 70 different countries.
[51]	2022	✘	✘	✔	Covert communication techniques.
[48]	2016	✘	✔	✔	Gap between practice and research.
[53]	2016	✔	✘	✔	Comparison of censorship resistance systems.
[54]	2015	✔	✔	✘	Censorship detection systems & platforms.
This	2025	✔	✔	✔	Measurements & circumvention techniques.

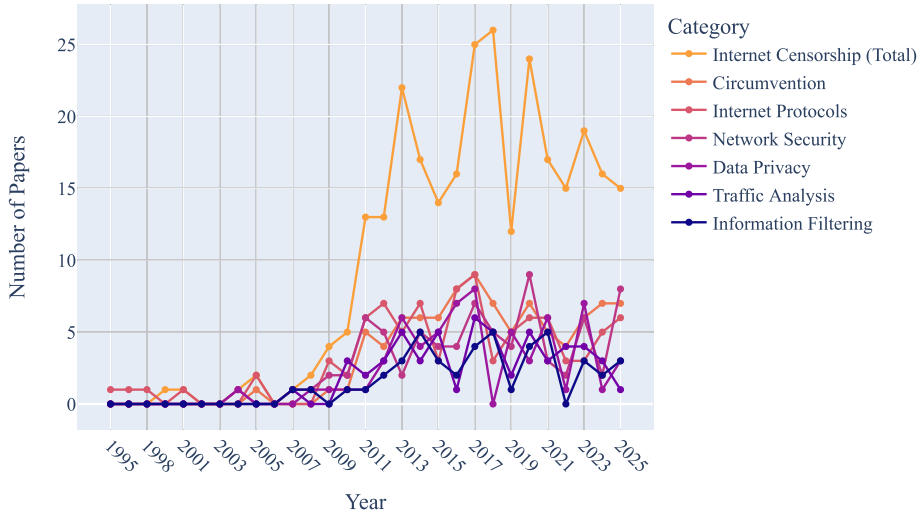

Fig. 2. Number of papers by research category related to Internet censorship from 1995 to 2025, based on Scopus data. All categories include only papers that directly pertain to Internet censorship.

as concerns about circumvention system stability due to volunteer recruitment and reliance on commercial participants that could become security/privacy risks.

One of the first surveys on Internet censorship detection is presented in [54]. It analyzes and discusses various censorship detection methods and architectures, and reviews existing tools and platforms for censorship detection. The survey also proposes a characterization scheme to compare different detection methods. The authors find that censorship detection systems function similarly to large-scale network measurement platforms; for instance, they face common challenges such as ensuring global measurement coherence and managing large network traffic loads. However, they also face unique issues due to adversarial interference from censors, who may block or tamper with probes and communications or even run their own rogue probes in crowd-sourced circumvention systems.

The authors of [52] present a contemporary survey of methodologies, trends, and challenges in Internet censorship measurement research. They begin by describing fundamental censorship techniques, including filtering based on IP, TCP, UDP, DNS, and TLS protocols. The survey then explores how interference with each of these protocols can be measured. In addition, the authors analyze and compare large-scale platforms used for measuring Internet censorship, and conclude with a discussion on both technical and societal trends and challenges associated with censorship.

Unlike previous surveys, our paper looks at Internet censorship measurement and circumvention from a different angle. While we do provide essential background and fundamentals, we also address aspects that are often overlooked, such as the interdependence between censorship measurement and circumvention research. We focus primarily on recent work from the past ten years, aiming to present the most up-to-date state of the field. We include a summary of the ethical challenges researchers face in this area. To give a more complete picture, we examine the geographical focus of existing studies and point out both well-studied and potentially understudied regions. Additionally, we examine the prevalence of longitudinal studies, assess the usage of various network protocols in measurement and circumvention research, and introduce a more detailed taxonomy of circumvention techniques.

3.1. Analysis of the evolution of literature

In this section, we present a set of visualizations that illustrate the evolution of research on Internet censorship over time, as well as the trends and tendencies within different research subareas. These figures were generated using the method introduced in [55]. The methodology, described in Section 4.3, employs Natural Language Processing (NLP) techniques and agglomerative clustering to categorize research themes based on keywords extracted from the entire literature on Internet censorship.

3.1.1. General growth of internet censorship research

Over the past decade, academic research on Internet censorship has grown considerably. Initially, studies were sporadic and focused on specific case studies, such as isolated censorship incidents or country-specific analyses. However, with a growing number of countries implementing restrictive Internet access policies, research has expanded significantly.

Fig. 1 depicts the growth in the number of research papers related to Internet censorship from 1995 to 2025. This figure reveals an increase in publications, with significant growth observed after 2008. This trend suggests growing academic interest in Internet censorship and increased government intervention in digital spaces.

3.1.2. Thematic trends in internet censorship research

The study of Internet censorship encompasses a range of topics, from technical assessments of censorship mechanisms to the social and political consequences of restricted online access. Research has been categorized into several themes that reflect different aspects of censorship. These themes include censorship circumvention, Internet protocols, network security, data privacy, traffic analysis, and information filtering.

Figs. 1 and 2 visualize the breakdown of research themes within the field. The themes include general Internet censorship research as well as closely related subfields such as network security and traffic analysis, considered only in the context of censorship-focused studies. This means, for example, that Fig. 2 does not cover all work on traffic analysis, but specifically those papers addressing traffic analysis in the context of Internet censorship.

4. Methodology

4.1. Semi-systematic literature review

Since Internet censorship research is a vast field, this literature review focuses predominantly on the last decade (2015–2025) of academic research on Internet censorship measurement and circumvention techniques. The paper follows the methodology of a semi-systematic literature review as defined in [56,57]. The semi-systematic literature review

is well-suited for synthesizing research in fields where methodologies vary, such as Internet censorship research. Unlike fully systematic reviews, which emphasize exhaustive and often quantitative analyses [57], a semi-systematic review focuses on identifying key themes, trends, and research gaps within a defined scope [56].

The first part of this review briefly summarizes existing review articles and can therefore be categorized as a tertiary literature review. The remainder of the paper follows the semi-systematic literature review methodology, reviewing primary studies with a specific focus on those published recently. The review process involved multiple stages, beginning with the identification of relevant databases and repositories. Each identified paper was counterchecked against the inclusion and exclusion criteria (see Section 4.2). Subsequently, key attributes of each study were recorded. For Internet censorship measurement studies, the key attributes are the used network protocols, the geographical focus, the timespan of the measurements, the methods, the primary findings, and the proposed tools or platforms. Internet censorship circumvention research was classified based on network protocols, the type of circumvention technique, network throughput, longitudinal or short-term nature of the study, the purpose of the technique (*e.g.,* purely for bootstrapping and/or for operational use), application-agnosticism, and real-world evaluation.

This classification enabled a thematic analysis of trends, highlighting advancements in measurement techniques (*e.g.,* the creation of large-scale measurement platforms), and visualizing the geographical focus of measurement studies.

4.2. Search strategy & inclusion and exclusion criteria

Initially, a comprehensive literature search was conducted, identifying relevant papers published in peer-reviewed journals and conferences. Relevant articles were retrieved from IEEE Xplore, the ACM Digital Library, Scopus, and the Google Scholar index, as well as curated open access archives such as CensorBib [58] and using backward snowballing [59].

To narrow the scope to high-quality and thematically relevant studies, inclusion and exclusion criteria were established. The criteria for inclusion were peer-reviewed articles that contain the keywords “Internet censorship” and have either conducted a study that *(i)* measures Internet censorship in the real-world or in a simulated environment, *(ii)* proposes a novel measurement and/or circumvention technique or incremental improvements thereof, or *(iii)* applies existing security and/or privacy-enhancing techniques (*e.g.,* a steganographical method) in the context of Internet censorship circumvention.

Excluded articles are studies that *(i)* conducted Internet measurements not directly related to uncovering censorship, such as [60], studies that *(ii)* measured the detectability of circumvention techniques without proposing new tools or improvements, and *(iii)* “pure” steganography papers. Furthermore, *(iv)* studies were excluded if they dealt primarily with legal, political, or social dimensions of censorship without a strong technical component (*e.g.,* [61]), unless they provided contextual insights critical for understanding technical developments. Lastly, *(v)* articles that measure pure behavioral or keyword-based censorship were also excluded, such as removing social media posts from controlled platforms based on keyword detection or image content analysis, as investigated in [62].

Applying the aforementioned criteria resulted in a total of 146 studies for further analysis.

4.3. Literature analysis plots methodology

The data used to generate the plots in Figs. 1 and 2 were obtained by querying the Scopus abstract and citation database [63]. These queries were intended to provide an initial exploratory visualization of publication trends over time using automated keyword clustering, rather than to gather a complete set of all relevant studies. Two separate queries were executed to extract relevant research papers:

1. A general overview of Internet censorship research was obtained using the query “internet censorship” in “Article title, Abstract, or Keywords”, and limiting the results to papers written in English, resulting in 429 papers.
2. A more focused dataset targeting measurement and circumvention studies was obtained using the query (“internet censorship” and “measurement” in “Article title, Abstract, or Keywords”) OR (“internet censorship” and “circumvention” in “Article title, Abstract, or Keywords”), resulting in 95 papers.

The analysis was conducted using the methodology described in [55], which relies on Natural Language Processing (NLP) to identify thematic patterns in large collections of scientific literature by extracting and analyzing trends. In this methodology, the textual metadata associated with each paper is transformed into numerical representations using the transformer model `all-mpnet-base-v2` [64]. These vector embeddings capture the semantic relationships between terms and enable the system to measure the semantic similarity between different research topics. An agglomerative clustering algorithm from the `scikit-learn` library [65] was used to group semantically related keywords according to their distances in the embedding space. Each resulting cluster represents a topic or theme within the research domain.

5. Fundamentals

This section briefly discusses the fundamentals of Internet censorship measurement and circumvention research, starting with an overview of relevant terminology and fundamental technologies related to censorship circumvention, such as Onion Routing and I2P.

5.1. Terminology

This subsection briefly provides definitions for relevant terminology. In addition, a comprehensive list of acronyms can be found in Table B.5.

5.1.1. Censors and censorship

According to [66], a *censor* is defined as “*a person responsible for examining books, films, works of art, or communications, and deciding whether to prevent parts or the whole of them from being seen or made available to the public because they are considered to be offensive or harmful, or because they contain information that someone wishes to keep secret, often for political reasons*”. Within the realm of the Internet, a *censor* is oftentimes an entity (*e.g.,* a nation state) that enforces rules to prohibit users from freely accessing information for various reasons. The motivations behind censorship are often shaped by political, economic, or cultural agendas aimed at controlling public discourse and access to information.

Other motivations for censorship are presented in [67]. They include regulatory aspects, such as non-GDPR compliant websites that block users with EU-based IP addresses, news and TV channels that are not accessible from abroad, or blocking due to security reasons. These may or may not be forms of censorship and cannot easily be categorized as such without knowing the deeper motivation behind them [67].

5.1.2. Network interference

In the context of this paper, network interference means any intentional disruption or alteration of normal network traffic. This is often used as a more neutral term to describe censorship. In other contexts, network interference can refer to disruptions at the wireless signal level, for instance, frequency overlaps that lead to data collisions [68].

There are different forms of network interference. Censors can outright block traffic (*e.g.,* based on protocol and IP addresses) or they may selectively filter traffic based on the traffic’s structure or content (*e.g.,* based on keywords in cleartext traffic [69,70] or even based on a payload’s entropy [71]). Another form of network interference, which is not as apparent as blocking or filtering, is traffic throttling. Depending on the application, throttling can render a service practically unusable [24].

5.1.3. Collateral damage

A term that is often used in conjunction with Internet censorship is collateral damage. It is defined as "*damage done to something or harm done to someone that is not intended*" [72]. In the context of censorship, this refers to the unintended disruption of access to legitimate services or content when attempting to censor communications. For censors, there is always a trade-off between the effectiveness of blocking specific content and the risk of causing such unintended consequences. This means that censoring services often leads to overblocking, where access to unrelated but co-hosted or similarly routed content is also restricted. Only countries with entirely self-contained digital ecosystems can afford to block major service providers outright, by filtering entire IP address ranges or DNS names, without experiencing significant collateral damage.

5.2. Censorship measurements

Internet censorship measurement is the study of how access to online content is blocked or restricted, typically by governments, ISPs, or other actors. Early research in this area often used terms like Internet filtering or network interference measurements [73]. Over time, the field adopted clearer terminology, "Internet censorship measurement" or "censorship analysis", as methods became more targeted and the political context more central. The first academic studies mentioning the term "censorship measurement" appeared in 2011 [74] and 2012 [75]. The authors of [74] investigated how China filters Internet traffic by analyzing China's AS-level topology. They identified how and where the Great Firewall of China (GFW) operates and showed that while most filtering happens at the national border via major ASes, some filtering also takes place deeper within regional (provincial) networks.

An even earlier investigation into how China's GFW functions was presented in [76]. The authors analyzed how the GFW performs keyword-based censorship in TCP connections. It checks the TCP payloads for restricted terms and, in case there is a match, the GFW sends TCP RST packets to the endpoints of the TCP connection, which effectively closes the connection. However, the authors found that the TCP RST packets can simply be ignored by using an *iptables* rule, and the requested packet (*e.g.,* an HTTP response) can still be received by the endpoint.

Back in 2006, [76] found that the GFW operates entirely symmetrically, meaning it filters both incoming and outgoing traffic using the same content inspection rules. Nowadays, there is evidence that the GFW operates asymmetrically [31]. Additionally, more recent studies [18,31,77] show that the filtering mechanisms of today's GFW are far more advanced and less error-prone. Another notable early work in the field is [78]. The authors developed ConceptDoppler, a method that tests for censored keywords in HTTP traffic. Using ConceptDoppler, they studied the GFW's keyword-based censorship by actively probing the firewall from outside the country. They continuously tested for censorship across different regions to detect when and where specific keywords were added to the blacklist. To efficiently identify new potentially blocked terms, they applied Latent Semantic Analysis (LSA) to a Chinese Wikipedia corpus, and were able to identify 122 filtered keywords.

In order to measure censorship, one typically requires the following setup: a machine, controlled by the researcher, in the target network environment (*e.g.,* within the censored country or region). This machine acts as a vantage point from which the researcher can perform network measurements, such as attempting to resolve domain names and/or access websites. The measurements are then compared to those taken from an uncensored control location to identify differences. These differences can then point to potential instances of censorship.

Further details on different measurement methods and the ethical considerations can be found in Section 6. A dedicated description of large-scale Internet censorship measurement platforms can be found in Section 6.3.

5.3. Censorship circumvention

Censorship circumvention refers to a range of techniques and systems developed to bypass restrictions imposed on information access. These methods enable users to access blocked content. Trying to bypass restrictions is part of a bigger struggle between people who want to control information and those who want to access it freely. It is often compared to a "cat-and-mouse game" between censors and people trying to circumvent them [79].

Early research on censorship circumvention introduced Freenet [80], a decentralized, anonymous, censorship-resilient file sharing system. Infranet [81], introduced in 2002, is an early circumvention scheme that disguises requests as benign HTTP traffic and delivers content steganographically inside images.

Today, one of the most widely used circumvention systems is Tor (see Section 5.5), which is a volunteer-run Onion Routing system that provides anonymous communication by encrypting and routing traffic through a network of relays. As later shown, there are different types of circumvention techniques. Most fall into one of two categories: routing-based and obfuscation-based circumvention. All further contemporary censorship circumvention methods are surveyed in more detail in Section 7.

5.4. Bootstrapping

One of the central challenges in censorship circumvention is establishing a covert communication channel in the first place, before any actual communication can occur to bypass censorship. This initial step, sometimes referred to as signaling [82], registering [83], rendezvousing [84] or proxy assignment [85], is commonly known as bootstrapping. Bootstrapping involves acquiring information about the circumvention system, such as proxy addresses or cryptographic keys, over a communication channel that passes through censor-controlled infrastructure. This process must enable genuine users to obtain access efficiently, while also preventing censors from discovering and subsequently blacklisting the bootstrapping endpoints.

Some circumvention methods, like domain fronting [86–88], attempt to bypass the need for bootstrapping or explicit proxy assignment altogether. However, these approaches face other practical limitations and, therefore, are unsuitable in many scenarios. Such limitations include being dependent on the cooperation of large CDNs and being restricted to use certain protocols only (such as HTTPS). As a result, bootstrapping (in particular proxy assignment) remains a difficult problem for most widely deployed circumvention systems [85].

Current proxy assignment mechanisms often rely on heuristics and struggle to distinguish between genuine users and censoring agents. This makes them vulnerable to enumeration attacks, where censors systematically discover and block proxies [85]. To mitigate this, bootstrapping techniques increasingly rely on steganographic methods by hiding secret information within seemingly "normal" communication.

In Tor, for example, censors continuously try to enumerate and block bridges [89]. Tor bridge enumeration can be partly prevented by using pluggable transports like obfs4 [90], which obfuscate traffic to make it indistinguishable from other encrypted protocols. Frequent rotation of bridges and limiting their exposure (*e.g.,* not listing them publicly or sharing them more selectively) also helps to reduce the risk of discovery.

5.5. Onion routing

The concept of Onion Routing was first introduced in 1996 by David Goldschlag, Michael Reed, and Paul Syverson in [91]. It is a routing scheme developed to handle real-time bidirectional traffic without revealing the identity of the receiver to the sender, and vice versa. A bidirectional communication channel is established by selecting a route that consists of a series of nodes, where each node only knows its immediate predecessor and successor [92]. A message is protected with multiple layers of encryption, each layer corresponding to one of the

selected nodes in the route. This method is often compared to the layers of an onion, where the outermost layer is decrypted by the first node, revealing the address of the next node, and so on, until the message reaches its final destination. Each intermediary node decrypts only its layer, forwards the message to the next node, and never has access to the original message content or the full path [92].

In the years following their initial proposal of Onion Routing, the three original authors complemented their idea with a series of other papers [93–96], in which they show how Onion Routing complicates traffic analysis and present implementation-specific details such as access configurations to an Onion Routing network.

The most well known instantiation of an Onion Routing scheme is "The Onion Routing" (Tor) by Roger Dingledine, Nick Mathewson, and Paul Syverson [97]. This overlay network implementation, first openly deployed in 2002 [8], includes features such as perfect forward secrecy, TCP multiplexing, congestion control, and end-to-end integrity checking [92]. In 2006, the Tor Project non-profit organization was founded, which has since been maintaining and improving Tor [8]. Over the years, the organization has expanded its efforts beyond the core Tor software to include tools and services that promote online privacy and freedom of expression. Since then, several enhancements have been added to the original implementation. As of 2025, the Tor Browser includes support for multiple different *pluggable transports*, which are tools that help circumvent Tor censorship [8]. Among these, Snowflake [98,99] has gained popularity since it enables users to bypass censorship by using temporary proxy relays provided by volunteers (see Section 7.2.3). Other notable extensions and tools include meek [100–102], obfs4 bridges [90], and WebTunnel [103].

5.6. Garlic routing

Garlic Routing is an extension of Onion Routing designed to enhance anonymity and resistance to traffic analysis. The technique is named "Garlic Routing" because, much like a head of garlic consists of multiple cloves, a single garlic message contains multiple "cloves", or individual messages. Unlike traditional Onion Routing, which encrypts and sends one message at a time, Garlic Routing bundles several encrypted messages together into a single message. Each individual message within this garlic bundle can have its own set of delivery instructions [104]. The bundling of multiple messages and use of separate tunnels make it significantly more difficult for an external observer to analyze traffic patterns or deduce communication relationships. Since each bundle contains multiple messages, it becomes much harder to fingerprint the traffic.

Additionally, Garlic Routing uses unidirectional tunnels, meaning that requests and responses travel along separate paths [105]. In comparison, standard Onion Routing typically implements bidirectional communication on the same tunnel.

5.7. I2P

The Invisible Internet Project (I2P) is a fully encrypted, decentralized, peer-to-peer (P2P) overlay network designed to enable censorship-resistant, anonymous communication [9]. It originated from a fork of Freenet (now Hyphanet [10]), which is also a censorship-resistant peer-to-peer network.

The main differences between I2P and Tor are that Tor is a partially centralized, onion-routed network optimized for anonymous access to the public Internet, whereas I2P implements a type of Garlic Routing and is fully decentralized, with a distributed network database and peer selection. Tor has a large number of exit nodes and is optimized for exit traffic, while I2P is optimized for hidden services that can be accessed faster than in Tor [9].

6. Internet censorship measurement research

Ethical considerations are a critical foundation for any measurement study, whether or not it directly involves end users and their devices. This section begins with a discussion of the ethical concerns that arise and how ethical guidelines have evolved over the last decade. The main body of this section presents an analysis of various aspects of censorship measurement research, based on 76 studies presented in Table 2: First, the geographical focus of measurement studies in recent years is discussed. Secondly, the different study durations are analyzed, followed by a brief look at the different network protocols used to measure censorship. Lastly, the proportion of longitudinal studies is discussed. In recent years, large-scale measurement platforms have emerged as part of a big push against censorship. They provide worldwide coverage and different probing techniques for all widely-used network protocols. These are discussed in Section 6.3.

6.1. Ethical considerations

Internet censorship measurement research often involves engaging with sensitive political environments, violating local regulations (unintentionally or intentionally), and interacting with vulnerable populations. These challenges raise complex ethical questions about the risks to individuals – both to external researchers and those participating from within censored regions (*e.g.,* political activists). This section briefly summarizes how ethical considerations have been addressed by researchers over the last decade and what prompted the efforts to minimize risks to uninformed individuals.

According to [183], there are three main methods for collecting censorship measurements. The first *(i)* involves sending researchers to perform direct measurements, but this is limited by logistics and safety concerns. The second approach *(ii)* relies on local citizens or activists to install measurement software, allowing for potentially more continuous monitoring, although this can also be risky and does not always ensure ongoing coverage. The third method *(iii)* leverages existing software infrastructure to indirectly collect data, offering wider, continuous coverage; however, this method raises ethical issues since users may be unaware of their involvement. While methods *(i)* and *(ii)* are ethically safer, they often lack the necessary data scope for certain research goals. Later on in 2015, [184] extended their initial ethical considerations described in [183] with more concrete steps to reduce the risk to volunteers *(ii)* participating in censorship measurement campaigns. They proposed measurement methods that minimize the observable differences in censorship measurement traffic as compared to "normal" network traffic and validated the methods using, inter alia, a signature-based Intrusion Detection System (IDS). However, their considerations are strongly based on assumptions and controlled experiments, but lack quantitative evidence of effectiveness in real-world deployments.

One of the first measurement systems that adopted the co-opt approach *(iii)* was Encore [180], which prompts visitors of modified webpages to make cross-origin requests to a URL that Encore aims to test for censorship. This paper sparked some discussions around ethical guidelines (or lack thereof) in the field of censorship measurements. The SIGCOMM 2015 Program Committee published the article, but also included a statement addressing a controversy that had arisen over the proposed measurement method's lack of informed consent. The Program Committee did, therefore, not endorse the method proposed in the paper. The authors of [180] argued that informed consent is not always applicable. In the case of Encore, they considered it impractical and arduous for the users to explain technical intricacies such as cross-origin requests and web trackers, arguing that providing such explanations would be overly technical and difficult for average users to understand. Furthermore, the authors feared that this requirement would dramatically reduce the scale of their experiments, as many users might opt out if they fully understood the experimental design. Consequently, a reduction of participants would have significantly decreased the effectiveness of the measurement, thus undermining the research goal of accurately identifying instances of censorship. The authors of [185] provide a comprehensive analysis of the Encore case, concluding that from a legal standpoint, Encore's methodological approach complied with U.S. law;

Table 2
Chronological synthesis of relevant Internet censorship measurement research (2015-2025).

Work	Year	Protocols	Countries	Duration	Longitudinal	Methods	Results	Data	Tool
[106]	2025	IP (incl. IPv6), TCP	212 countries & territories	6	○	Designed a global censorship measurement system that can uncover censorship without the need for in-country VPs. Scanned /24 IPv4 and /48 IPv6 subnets to identify unresponsive IPs that are checked for interference.	Scanned IPv4 and IPv6 ranges across 9483 ASes and integrated Geneva [107], which can be used to discover new packet sequences that can trigger censorship.	●	Mint [108]
[109]	2025	TLS, DNS-over-[TLS, HTTPS]	China, Iran, Russia	4	●	Evaluated TLS Encrypted Client Hello (ECH) server support in the wild using the Tranco Top 1M list and ECH censorship by sending ECH messages through censored regions.	Identified Cloudflare as the only provider to support ECH, although not always advertised. Confirmed direct censorship of ECH in Russia and indirect censorship in China and Iran by blocking unencrypted and encrypted DNS.	●	N/A
[110]	2025	DNS, HTTP	Iran	2	○	Used an in-country VP and a reference server in Germany. Used Tranco top 1M domains and Citizen Lab global test list for DNS measurements and the Tranco top 100 list and the Citizen Lab test list for Iran for HTTP tests.	Discovered that 87% of DNS responses injected the same IP address, found HTTP censorship for all HTTP methods, and instances of overblocking. Censorship is implemented on border nodes.	●	N/A
[111]	2025	QUIC	China	3	●	Measured QUIC Initial packet filtering in the GFW using 7 in-country and 6 out-of-country VPs, a custom Cloudflare quiche QUIC client.	Found differences between the QUIC blocklists and existing TLS, DNS and HTTP blocklists. Uncovered that QUIC Initial packet decryption degrades performance of GFW. Proposed a set of circumvention techniques.	●	GFWReport [112]
[113]	2025	TCP, HTTP, HTTPS	73 countries & territories	–	○	Created a DPI measurement framework which selects and deploys probes. Compared how different DPI middleboxes handle these probes and aggregated the outcomes into behavioral fingerprints.	Conducted 3 million measurements on 33K targets and demonstrated that 20–40 discriminative probes can differentiate many DPI deployments.	●	dMAP [114]
[115]	2025	TCP, TLS	Russia	36	●	Measured outbound Internet traffic throttling for a set of sensitive URLs based on data provided by OONI [37]. Analyzed over 86M data records gathered over a three-year period.	Identified the ASes that introduce the most delays. Found that popular social media apps are either blocked or that the mean TLS handshake time has increased significantly over time.	●	N/A
[116]	2025	UDP, DNS, HTTP	Iran	2.5	●	Measured the entire IP space of Iran using out-of-country VPs exclusively. Reverse-engineered blocking rules and tested for DNS poisoning, HTTP blockpage injection and UDP traffic disruptions.	Found 6.8M IPs subjected to DNS poisoning and HTTP blockpage injection and 5.4M IPs affected by UDP-based censorship. Identified overblocking of domains.	●	IRBlock [117]
[118]	2025	UDP, TCP, TLS, STUN	Saudi Arabia, UAE	–	○	Tested 9 messaging apps for VoIP censorship in person from inside the countries without the use of VPs. Evaluated the existence of call setup censorship and selective filtering of STUN packets.	Found app-specific fields in STUN headers that are used for fingerprinting-based blocking. Censorship is bidirectional and affects audio and video calls. Plan to integrate STUN probing into OONI [37] and Censored Planet [38].	●	[119]
[120]	2025	TLS, HTTP	China	9	●	Investigated regional censorship on a province level in China and its differences to GFW-based filtering. Used the Tranco top 1M and CZDS 227M domains.	Uncovered provincial-level HTTP Host field and TLS SNI censorship in Henan. Found differences to the GFW, *e.g.,* in that regional middleboxes only censor outgoing traffic but block 5x more domains than the GFW.	●	[121]

(*continued on next page*)

Table 2 *(continued)*

Work	Year	Protocols	Countries	Duration	Longitudinal	Methods	Results	Data	Tool
[122]	2025	IP, TCP, TLS, DNS, HTTP, SSH	Iran	20	●	Tested for censorship by packet injection, mid-connection dropping, IP blocking, protocol allowlisting. Experimented with server-side evasion, through TLS Client Hello injections and detection-resistant protocols.	Revealed decentralized censorship in Iran with differences across ISP and successfully deployed new methods to circumvent proxy censorship.	○	Psiphon [123]
[124]	2024	DNS, HTTP, HTTPS	21 countries & territories	–	●	Proposed new route-stable measurement methods that prioritize consistency. Conducted outside-in censorship measurements, including censorship-and-path aware traceroutes.	Investigated how ECMP routing changes observed censorship across strategies, protocols & countries. ECMP routing was found to result in censorship changes in 42% of IPs and 51% of ASes.	○	Monocle N/A
[31]	2024	HTTP, HTTPS	China	20	●	Bidirectional, continuous measurements of China's GFW; cross-protocol comparison to GFWatch's [125] DNS measurements.	Tested 1.02 billion FQDNs, of which 943K and 55K pay-level domains are censored; they reveal that the GFW's censorship is not symmetrical.	●	GFWeb [126]
[127]	2024	DNS	China	< 1	○	Exploited an out-of-bounds read vulnerability exfiltrated memory parts from the GFW using malformed DNS requests. Collected 13B leaks.	Extracted sensitive information by overflowing first or last label in DNS queries. Found 3233 domains triggering this vulnerability; which is now patched.	○	N/A
[128]	2024	TCP, TLS, DNS	India	–	○	Analyzed analyzed the accessibility of 220 Chinese mobile apps within India by investigating ISP-level filtering, geoblocking and SIM-card based-blocking.	Found no ISP-level censorship; but rather app publisher censorship using source IP-based geoblocking and/or SIM card-based blocking.	○	N/A
[129]	2024	IP (VoWiFi), DNS	219 countries & territories	1	○	Proposed a method to detect VoWiFi deployments and geoblocking, by mapping DNS records and probing ePDG servers.	Conducted 8555 domain and 47,902 IP discovery scans and found IP-based blocking on the DNS and predominantly on the IKE layer in Eurasia.	●	scanywhere [130]
[32]	2024	HTTP, HTTPS, DNS, Echo, Discard	223 countries & territories	48	●	Employed an decision tree-based unsupervised learning method to find domains with similar blocking patterns using 70 billion CensoredPlanet [38] data points.	Identified 15,360 HTTP(S) and 1166 DNS event clusters in 192 countries and 77 countries, respectively; found 100 ASes in 32 countries with persistent ISP blocking.	●	CenDTect N/A
[43]	2023	IP, HTTP, DNS	China, India, Kazakhstan	–	○	Developed a flow-level packet-manipulation detector and performed detection-resistant probing from multiple VPs in China, one VP in India and one in Kazakhstan.	Effective at shielding probing traffic from censors against China's GFW (up to 98%), and 100% in India and Kazakhstan. Demonstrated that evasion tools like Geneva [45] can be blocked by censors.	○	DeResistor [131]
[132]	2023	HTTP	221 countries & territories	–	○	Evaluated a supervised learning-based classification model that learns latent features from network traffic using Censored Planet [38] Quack datasets.	Created an end-to-end network-based censorship detection pipeline which identified previously unknown censorship block page responses.	●	CPLearning [133]

(*continued on next page*)

Table 2 *(continued)*

Work	Year	Protocols	Countries	Duration	Longitudinal	Methods	Results	Data	Tool
[18]	2023	TLS, HTTP, SSH, SMTP, FTP, DNS-over-TCP	China	10	●	Conducted thorough measurements of China's GFW's encrypted traffic censorship system, *e.g.,* tested if the fraction of bits set in payloads and the position of ASCII characters trigger the system.	Reverse-engineered the blocking rules used by the Chinese GFW to block encrypted traffic and proposed circumvention techniques, such as customizable payload prefixes.	●	[134]
[13]	2023	TCP, HTTP, TLS, DNS	Russia	12	●	Measured geoblocking in Russia after the invasion of Ukraine using OONI [37], Censored Planet [38] and other public data, and in-country as well as out-of-country VPs.	Identified increase in censorship, evidence of BGP withdrawals, and evidence of geoblocking for user abroad trying to access in-country domains. Observed emergence of a domestic CA.	●	GeoInspector [135]
[27]	2023	HTTP, HTTPS, DNS	India	–	○	Developed a mobile app to analyze censorship of 10,372 sites across 71 ASes from 25 Indian states. Collected a total of 8.9 million measurements.	Found substantial differences in blocking behavior across ISPs; found that widespread blocking occurs without any legal basis.	●	CensorWatch [136]
[137]	2023	HTTP, TLS, DNS	27 EU countries	10	●	Analyzed 1M OONI [37] measurements from 888 distinct ASes for the existence of blockpages, fake DNS responses or incorrect AS information.	Found 51 unique blockpages from 18 countries across 47 ASes. All countries maintain a gambling site blocklist and almost all a copyright blocklist.	●	OONI [37]
[33]	2023	ICMP, IP, BGP	155 countries & territories	48	●	Used a combined dataset on planned and spontaneous Internet outages using IODA [138], Access Now [139], and other sources to identify predictors for shutdowns.	Identified correlation between Internet shutdowns and days of elections, protests, and coups. Found that the fingerprint of spontaneous outages differs from planned ones.	●	[140]
[141]	2023	DNS	China, USA	7	●	Used 4.7M DNS censorship records from OONI [37] and Censored Planet [38] and trained (un)supervised models to detect DNS-based censorship.	Showed that (un)supervised models can detect new instances of censorship; both supervised and unsupervised models achieved high TPRs.	●	[142]
[143]	2023	HTTP	Russia	–	○	Investigated transit censorship by scanning IP address spaces of 18 neighboring countries. Searched for injected blockpage responses.	Revealed that 9/18 countries are affected by 7 ASes that tamper with transit traffic; localized censorship devices using TTL-limited forbidden GET requests.	○	N/A
[144]	2023	TCP, HTTP, HTTPS	247 countries & territories	< 1	●	Conducted a global study on connection tampering using Cloudflare CDN traffic. Uniformly sampled one of every 10,000 connections to a Cloudflare server for analysis.	Collected and analyzed connection tampering signatures and found 19 unique fingerprints. Assessed the scope of current test lists, exposing overlooked domains affected by tampering.	○	N/A
[42]	2023	DNS, HTTP, HTTPS	Turkmenistan	2	●	Used out-of-country probing to test 15.5M domains and reverse-engineered censored domain regex rules; further used Geneva [45] to discover new censorship evasion strategies.	Found >122K censored domains and different blocklists for every protocol; identified 6000 over-blocking rules that cause accidental filtering of 5.4M domains.	●	TMC [145]

(*continued on next page*)

Table 2 *(continued)*

Work	Year	Protocols	Countries	Duration	Longitudinal	Methods	Results	Data	Tool
[146]	2022	HTTPS	26 countries & territories	1	●	Conducted large-scale study of mobile app blocking using semi-automated measurements from VPS vantage points. Analyzed 5684 popular apps.	Found 3672 geoblocked apps, with Iran and Tunisia blocking the most apps (2256 and 2681). Identified many developer-blocked apps.	●	geodiff-app [147]
[28]	2022	DNS	China	–	○	Explored the impact of packet headers (local source IP and port) on DNS censorship on a large-scale. Employed a fuzzing strategy to vary ports and IPs. Used techniques from [38, 148], and [149].	Found a change in blocking behaviour in 37% of IPs across 56% ASes, when varying source ports, IPs, and TTLs. Identified the routes taken by DNS packets and the method of blocking.	○	BreadCrumb N/A
[150]	2022	IP, HTTPS, QUIC	Russia	–	●	Conducted in-country and remote measurements aimed at identifying where and how blocking occurs. Inter alia, created custom TLS ClientHellos to test SNI censorship.	Identified >1M endpoints behind the censorship system (TSPU) affected by IP-, SNI-, QUIC-based in-path blocking (outgoing direction); Found TSPUs deployed close to end user devices.	○	N/A
[151]	2022	DNS, TLS, HTTPS	85 countries & territories	6	●	Performed 315K measurements from >20K vantage points on 1.6K domains in 878 ASes; measured prevalence of encrypted SNI and DNS-over-HTTPS adoption.	Identified multiple countries, including China, Russia, and Saudi Arabia, that block encrypted domain names. Showed that ESNI adoption in TLDs is low.	●	DNEye N/A
[23]	2022	HTTP, HTTPS	Azerbaijan, Belarus, Kazakhstan, Russia	1	○	Used TTL-probing techniques to identify the location of censorship devices; applied HTTP(S) static fuzzing strategies to study the blocking behaviors.	Identified location and type of terminating hops; Showed that HTTP hostname padding, alternate HTTP method fuzzing, and TLS Client Hello SNI field fuzzing are successful.	●	CenTrace [152], CenFuzz [153]
[14]	2021	DNS, HTTP, HTTPS	Spain	47	●	Analyzed >3M OONI [37] data points focusing on websites advocating for civil rights, encrypted messengers, and information hubs related to the Catalan referendum.	Found evidence of advanced network interference techniques deployed by major ISPs; detected 16 unique blockpages, 2 DPI vendors, and 78 blocked websites.	●	N/A
[154]	2021	HTTP/3 over QUIC, HTTPS	China, Iran, India, Kazakhstan	2.5	○	Measurement of differences in domain availability across HTTPS and HTTP/3 over QUIC using in-country vantage points.	China, India and Iran employed HTTP/3 censoring during the measurement period, albeit less strictly than HTTPS censorship.	●	OONI [37] Extension
[77]	2021	DNS	China	9	●	Large-scale measurements of DNS censorship by China's GFW; identification of IPv4/IPv6 addresses of forged DNS responses.	Probed 534M distinct domains, 311K turned out to be censored; found a small number of websites that are unintentionally censored due to greedy regex.	●	GFWatch [125]
[44]	2021	HTTPS	China	–	○	Used censored in-country and uncensored out-of-country VPs and Geneva [107] to find 4 new evasion strategies. Measured how two GFW middleboxes behave differently.	Found evidence for a second HTTPS censorship middlebox as part of the GFW due to unique per-box TCP-layer bugs that inject TCP resets differently.	●	Geneva [107]

(*continued on next page*)

Table 2 *(continued)*

Work	Year	Protocols	Countries	Duration	Longitudinal	Methods	Results	Data	Tool
[69]	2021	HTTP, HTTPS	China	1	●	Performed keyword-based filtering tests using in-country and out-of-country VPS. Tested Chinese (simp. and trad.) and English keywords from four lists of censored terms.	Found TLS SNI-based censorship, and a change in forbidden keywords since 2014. Identified discrepancies between the GFW's Chinese chat clients' keyword blocklists.	●	N/A
[24]	2021	Echo, HTTP, TLS	Russia	2	●	Measured the characteristics and location of Twitter traffic throttlers using in-country vantage points and public data covering 401 unique Russian ASes.	Found that the throttler acts on Twitter domains and reduces the up- and downstream to 130–150 kbps; found consistency across ISPs	●	Censored Planet [38]
[155]	2021	DNS-over-[TLS, HTTPS]	Kazakhstan, Iran, China	1	●	Integrated DNSCheck into OONI [37] and tested the reachability of 123 DoT/DoH services using OONI's volunteer network.	Identified 50% of DoT endpoints blocked in Iran, predominantly Cloudflare's and Google's; no blocking was observed in other ASes.	●	DNSCheck [37]
[15]	2021	IP, DNS, HTTP, BGP	Myanmar	3	●	Analyzed IODA [138], OONI [37] and proprietary datasets during and after the military coup in Myanmar (Feb. 1 2021).	Observed outages on the coup day, regular nightly outages and restriction in cellular connectivity as well as dedicated TCP/IP and DNS blocking.	●	[37,138]
[2]	2020	IP, HTTP, HTTPS, DNS, Echo, Discard	221 countries & territories	20	●	Utilized six protocols and geographically diverse in-country vantage points; performed measurements via TCP side channels, echo servers, web servers.	Collected and analyzed 21 billion data points from 95,000 vantage points (at time of publication) and found censorship activities in 100 + countries.	●	Censored Planet [38]
[71]	2020	HTTP, HTTPS	China	4	●	Investigated Shadowsocks censorship by using Shadowsocks-libev clients on five in-country VPSes and the resp. servers on five out-of-country VPSes.	Identified that the GFW analyzes the initial Shadowsocks packet using its entropy and size, and subsequently probes suspected Shadowsocks servers. Presented a mitigation strategy.	●	[156]
[16]	2020	HTTP, HTTPS	Kazakhstan	–	●	Employed in-country and remote measurements of popular domains to identify censorship. Located interception points using a TTL-based technique.	Identified a 21-day period in 2019, during which Kazakhstan partially performed a HTTPS Interception attack on 37 unique domains.	●	Censored Planet [38]
[1]	2020	Echo, HTTP, HTTPS	>190 countries & territories	3	●	Developed a framework to monitor blockpage-based content filtering. Analyzed 379M measurements from 45,000 VPs. Clustered blockpages based on HTML content and visual similarity.	Detected 90 blockpage clusters deployed in 103 countries within firewalls, ISPs, commercial products. Found deployments in 36 out of 48 "not free" or "partly free" countries.	●	FilterMap [157]
[70]	2020	TCP, DNS	Russia	–	○	Probed from residential and non-residential VPs spanning 408 unique ASes. Acquired blocklists from Russian activists and checked for TCP/IP, keyword-based blocking and DNS manipulation.	Identified higher levels of censorship and more blockpage injections in residential vantage points. Varied censorship across ISPs complicates measurement and resistance.	●	Censored Planet [38]

(*continued on next page*)

Table 2 *(continued)*

Work	Year	Protocols	Countries	Duration	Longitudinal	Methods	Results	Data	Tool
[36]	2020	DNS, HTTP, ICMP	62 countries & territories	21	●	Leverages commercial VPN servers as vantage points to detect DNS manipulation, TCP packet injection and block pages.	ICLab detected over 3500 unique censored URLs through 53,906,532 measurements of 45,565 unique URLs and 234 ASes.	●	ICLab [39]
[158]	2020	DNS	China	9	●	Sent 2.8 billion DNS queries between out-of-country and in-country VPS, based on the Alexa top 1M domain test list; fingerprinted different DNS injectors.	Observed a total of 119.6 million forged DNS responses and found that one DNS query can trigger several spoofed responses. Found that one injector reuses the IP TTL from the probes.	●	Triplet Censors [159]
[19]	2020	TCP, DNS, HTTP, HTTPS	India	–	○	Investigated techniques used by six different Indian ISPs; proposed an automated HTTP censorship detection method.	Created a collection of potentially blocked websites and identified inconsistencies in blocklists across ISPs.	○	N/A
[160]	2020	TCP, DNS, HTTP, TLS	Saudi Arabia	26	●	Used six in-country VPs routing over the three major ISPs to test the availability of 18 social media apps and the Alexa's top 500 websites list.	Observed a decrease in HTTP filtering in the categories *Adult*, *Shopping*, and *Games* and a decrease in HTTPS filtering in the *Shopping* category over the entire measurement period.	○	N/A
[161]	2020	DNS, HTTP, HTTPS	Iran	–	○	Used Geneva [107] to identify evasion strategies and six in-country residential and non-residential VPs to perform measurements using the Alexa top-20,000 websites list.	Found that the protocol filter only operates uni-directionally (outgoing) on TCP ports 53, 80, and 443. Discovered four evasion techniques that successfully circumvented the filter.	○	Geneva [107]
[45]	2019	TCP, HTTP, HTTPS	China, India, Kazakhstan	–	○	Conducted in-lab and real-world tests using a genetic algorithm-based method that can perform mutations and crossovers on TCP traffic to identify new evasion strategies.	Could re-derive existing and new evasion strategies for the GFW of China, ISP censorship in India, and HTTPS interception in Kazakhstan and showed that some existing strategies are outdated.	●	Geneva [107]
[162]	2019	TCP, DNS, TLS	14 countries & territories	–	○	Studied TLS SNI- and ESNI-based censorship. Used Alexa top 1 million websites, one in-country, and one out-of-country VP.	Found 10.9% of websites in the Alexa top 1 million list supporting ESNIs. Did not identify dedicated ESNI blocking behaviour.	○	N/A
[163]	2019	SSH	Iran	6	●	Presented a comparative analysis of Psiphon [123] censorship during the Iranian elections of 2013 and 2016 based on analytics data from Psiphon.	Discovered that the Iranian regime refined its censorship tactics by utilizing more targeted blocking and by strategically timing its actions.	●	N/A
[3]	2019	HTTP, HTTPS, DNS	164 countries & territories	1	○	Measured I2P anonymity network censorship using VPN servers run by volunteers. Conducted 54K measurements from 1.7K different locations.	Found I2P censorship in China, Iran, Oman, Qatar, and Kuwait. The detected censorship techniques included DNS-based and TLS SNI-based blocking.	○	N/A

(*continued on next page*)

Table 2 *(continued)*

Work	Year	Protocols	Countries	Duration	Longitudinal	Methods	Results	Data	Tool
[164]	2018	TLS (Tor)	China	–	○	Used a VP in China, Tor bridge relays in the US, Canada, and the UK, and the TCIS [165] to analyze Tor bridge blocking.	Found 934 unique scanner IPs used by the GFW. Concluded that ignoring the scan requests is effective in preventing relays from being censored.	●	[166]
[167]	2018	Echo, HTTP, HTTPS	40 countries & territories	–	○	Introduced an application-layer blocking measurement method. Tested 4458 ASes, 100,000 domains from >100 vantage points (echo servers).	Found higher rates of censorship for some domains in 7 countries; found less censorship when using the Discard Protocol.	●	Quack [168]
[169]	2018	IP, DNS, HTTP, HTTPS	Pakistan	8	●	Consolidated censorship measurements and circumvention in one tool; upon detection of censorship, dynamically chooses a circumvention method with a short page load time.	Detected blockage of Twitter and Instagram between November 25–28, 2017 in Pakistan; conducted a case study which improved page load times via TOR or Lantern.	○	C-Saw N/A
[25]	2018	TCP, DNS, HTTP	India	18	●	Used heuristics to identify different types (and locations) of censorship in nine ISPs. Sent probes from their own machines, 50 planetlab hosts, ten cloud-hosted VMs and ten institutional machines.	Demonstrated inaccuracies of OONI [37]; Found differences in DNS and HTTP censorship across ISPs and tested circumvention techniques by fuzzing HTTP GET requests.	○	N/A
[20]	2018	HTTP	177 countries & territories	–	○	Tested for CDN geoblocking using thousands of residential VPs (Luminati proxy service) and the Alexa Top 10K sites.	Identified seven CDNs that enable geoblocking and found the highest rates of geoblocking in Iran, Syria, Sudan, Cuba, and Russia.	○	Lumscan N/A
[170]	2017	HTTP, HTTPS, DNS, TLS (TOR)	Cyprus	3	○	Performed measurements on 5 ISPs using a custom volunteer-based OONI [37] probe and used publicly available data from OONI reports.	Found multiple, previously unknown, instances of DNS hijacking-based censorship in 3 ISPs and identified several block pages.	●	bet2512 [171]
[26]	2017	HTTP, HTTPS	China, Iran	–	○	Automated the detection of middlebox filtering policies and their on-path position. Used a testbed and real-world measurements to test identify and circumvent DPI-based classifiers.	Found that classifiers rely on keywords in HTTP payloads, TLS SNI and other protocol-specific fields, and inferred viable circumvention strategies.	●	lib·erate [172]
[173]	2017	DNS	China	–	○	Developed a framework for automatic detection of filtered URLs by analyzing linguistic patterns in already filtered web pages.	Uncovered 1355 poisoned domains and proposed a metric to assess how effectively the language patterns of a base domain can be used to identify additional blocked domains.	○	FilteredWeb N/A
[174]	2017	TCP	~ 180 countries & territories	< 1	●	Developed a robust TCP side channel measurement method with sequential hypothesis testing that improves upon the Hybrid-Idle Scan.	Found inbound/bidirectional disruption occurring more often than outbound-only filtering; found high levels of filtering in China, Iran, Sudan.	○	Augur N/A
[175]	2017	DNS, HTTP, HTTPS	151 countries & territories	–	●	Conducted 13,594,683 DNS manipulation measurements using 6020 resolvers and tested 2303 unique domains.	Developed consistent, verifiable, ethical detection techniques which detected 41,778 manipulated responses, in 58 countries across 1408 domains.	○	Iris [149]

(*continued on next page*)

Table 2 *(continued)*

Work	Year	Protocols	Countries	Duration	Longitudinal	Methods	Results	Data	Tool
[176]	2017	ICMP, HTTP, BGP	219 countries & territories	12	●	Combined ICLab [39] with boolean network tomography to identify censoring ASes on-path by analyzing BGP churn. Used 4.9M measurements from VPs located in 539 ASes.	Identified 108 censoring ASes, 32 of which impose censorship affecting users outside their own jurisdiction due to leakage of censorship policies through BGP.	●	N/A
[29]	2017	DNS, HTTP	USA	–	○	Automated HTTP GET request fuzzing; used bidirectional probing to fingerprint the behavior of censorship devices.	Uncovered 76 web filters in the New York City metropolitan area that implement different blocklists and identified viable circumvention paths.	○	Autosonda N/A
[177]	2017	TCP, DNS	India	–	○	Mapped key AS and router-level paths and measured how frequently paths to popular sites pass through certain ASes. Probed IP prefixes to identify ASes for DNS injection.	Identified differences in censorship enforcement across ISPs. Found that 10 ASes intercept >95% of paths and 20% of ASes can interrupt connectivity for all Indian ASes.	○	N/A
[46]	2017	TCP	China	2	●	Performed measurements of TCP-level evasion on the GFW using 11 in-country and out-of-country vantage points and the Alexa Top websites list.	Translated their measurement results into new evasion strategies for HTTP, DNS, VPN, and Tor, that leverage TCP Control Block manipulation.	○	INTANG [178]
[21]	2016	DNS	China	–	○	Sent DNS requests to 1871 public DNS servers in China trying to resolve 15 commonly censored test domains.	Spotted that requests received both a valid and a poisoned response from the GFW. Identified nine IP addresses commonly used by the GFW.	○	N/A
[179]	2016	DNS, HTTP	169 countries & territories	24	●	Analyzed content distribution networks and network interference in 10,000 popular domains using a single external vantage point.	Detected 4819 instances of ISP-level DNS hijacking in 117 countries, showing growing interference of domain resolutions.	○	Satellite [148,149]
[30]	2016	TCP, DNS, HTTP	Pakistan	6	●	Compared censorship behaviors in five ISPs using a custom version of OpenWRT and performed up to six probing tests per day and VP.	Found that WiTribe, PTCL, and Nayatel ISPs perform DNS tampering; Wateen and Qubee employ use HTTP tampering.	○	N/A
[180]	2015	TCP, HTTP, DNS	170 countries & territories	7	●	Prompted visitors of modified webpages to make cross-origin requests to a URL that Encore aims to test for censorship.	Encore detects if URLs are filtered, but not how. Conducted 141,626 measurements from 88,260 IPs and confirmed well-known censorship cases.	○	Encore N/A
[181]	2015	TCP	China	1	●	Introduced a side-channel analysis technique using the Linux kernel's SYN backlog to distinguish between TCP SYN and SYN/ACK dropping.	Observed that the GFW blocks SYN/ACK segments entering China; side-channel method can measure the reachability of TOR nodes.	○	N/A
[34]	2015	TCP, HTTP, DNS	77 countries & territories	60	●	Analyzed data from the OpenNet Initiative. Analyzed measurements from over 90K distinct URLs, involving 286 distinct ISPs.	Identified blocking methods in understudied countries, and found significant variability in censorship deployments across regions.	●	N/A
[182]	2015	TCP, HTTP, DNS	31 countries & territories	4	●	Implemented a user-based censorship analysis platform, that consists of >200 distributed probes and tested 16K different targets.	Identified instances of censorship in Italy, Pakistan, and South Korea. Published tool as lightweight Linux client for voluntary participation.	○	UBICA N/A
[22]	2015	HTTP, DNS	Greece	3	○	Investigated gambling websites censorship using the blacklist from the Hellenic Gaming Commission. Collected data from eight ISPs using "ooniprobe".	Found under- and overblocking and instances of DNS manipulation and packet payload inspection performed by ISPs.	●	OONI [37]

however, global measurements leading to violations of local laws could not be ruled out.

The discussion revolving around Encore showed that there had been a lack of widely-accepted ethical standards for this research field. The authors of [180,183,185] arguably laid the groundwork for developing ethical norms in this research area by initiating discussions with ethics experts. Since then, large-scale Internet censorship measurement research adheres to noticeably more stringent ethical standards [2,18,25,31,45,146]. Measurement probes now commonly include a notice, informing affected parties about the possibility to opt-out by sending a simple response to the researchers. Another important aspect which seemingly all anti-censorship researchers agree on, is reducing the risk to uninformed natural persons. Researchers tend to prefer using enterprise-level infrastructure to conduct measurements (*e.g.,* VPN and Cloud providers) and try to avoid using privately-owned vantage points wherever possible, especially in cases where persons did not actively register to take part in a measurement study. Conversely, when involving natural persons on a voluntary basis, the principle of informed consent should be followed. Informed consent involves researchers clearly explaining a project's details and its risks to participants, who then decide whether to participate or withdraw. It is required when research could harm individuals, especially if it involves private information or information associated with natural persons, such as endpoint IP addresses in censorship measurement studies [186].

Studies, such as [124,167], follow the ethical principles and guidelines presented in the Belmont Report [187] and the Menlo Report [186]. The Belmont Report summarizes fundamental ethical principles to guide research involving human subjects and formulates guidelines to ensure compliance with these principles. The report stresses the importance of informed consent, an assessment of risks and benefits, as well as the careful selection of subjects [187]. The Menlo Report builds upon the Belmont Report and introduces ethical guidelines for Information and Communication Technology (ICT) research. Its aim is to reinterpret existing ethical guidelines and apply them to ICT studies that interact, directly or indirectly, with human subjects.

For example to ensure maximum compliance with ethical guidelines in ICT research, [167] leveraged globally distributed echo servers as vantage points, completely avoiding the need to involve volunteers, [16] used echo servers, web servers, and side-channel measurements, [36,158] performed their measurements using commercial in-country and out-of-country Vantage Points (VP), whereas [70] relied on a combination of residential and non-residential VPs. Researchers are also minimizing the number of per-VP measurements by distributing the load across different VPs [70]. OONI [37] mostly relies on a network of volunteers who install OONI probes locally. Despite its reliance on volunteers' personal infrastructure, OONI explicitly informs the user about the potential risks involved with running an OONI Probe desktop app during the onboarding stage.

In conclusion, the field of Internet censorship measurement research has evolved significantly in response to the ethical challenges posed by collecting data in politically sensitive environments and involving human subjects. Early debates, such as those sparked by the Encore system in 2015, highlighted the lack of clear ethical guidelines, particularly those concerning potential risks to uninformed individuals. Since then, researchers have worked to address these concerns by adopting more rigorous ethical standards, drawing on frameworks like the Belmont [187] and Menlo [186] Reports. Contemporary studies heavily rely on enterprise-level infrastructure, informed consent from volunteers, and transparency in communicating risks. As the field continues to grow, ethical considerations remain central to protecting participants.

6.2. Analysis

In total, 76 Internet censorship measurement studies (see Table 2) were analyzed with respect to the following aspects:

- *Protocols* refers to the Internet protocols used for measuring censorship between a vantage point and a potentially blocked endpoint.
- *Countries* lists the number of countries and territories that were involved in the experiments, not just the number of countries that were found conducting censorship.
- *Duration* specifies the study duration in months. Studies in which the duration was not explicitly mentioned are denoted with "–".
- *Longitudinal* describes the nature of the study. Longitudinal Internet censorship measurement studies continuously repeat their measurements using the same vantage points and (potentially censored) endpoints over longer periods of time.
- *Data* refers to whether the data published in the original paper was still accessible at the time this publication was written (2025).
- *Tool* points to the implementation of the measurement technique. It has been marked as "N/A" if there was no link provided to an implementation, and if no GitHub repository or entries on other open-source platforms, such as *paperswithcode.com*, were available.

6.2.1. Geographical focus of measurement studies

Fig. 3 illustrates the geographical focus of the Internet censorship measurement research, which is based on 65 studies that explicitly specified the countries investigated. It is essential to emphasize that this map does not represent the actual prevalence of censorship in specific countries or territories. Instead, it reflects the areas prioritized by researchers for investigation; more specifically, it shows the number of scientific studies that have measured censorship in a certain area. It is evident that the geographical focus of Internet censorship measurement research in the last decade lies on China, India, Iran, and Russia, but also the United States, United Kingdom, Canada, Australia, Kazakhstan, Pakistan, Turkey, Spain, France, Hungary, the Netherlands, Saudi Arabia, the United Arab Emirates, Indonesia, South Korea, and Japan have been notable subjects of investigation. These countries have drawn research attention due to a variety of factors, including the implementation of distinct censorship practices, the geopolitical significance of their information control policies, and the availability of resources and access for researchers. In some cases, the research interest in a region aligns closely with the prevalence of censorship within that country (*e.g.,* China). However, in other instances, there is little correlation between research focus and the actual censorship levels (*e.g.,* Cuba) [188]. This discrepancy can be attributed to several factors. For example, the scale and technical diversity of Internet censorship mechanisms in China make it a compelling target for researchers. Such systems attract researchers to reverse-engineer their operation. Furthermore, a censorship environment, that constantly changes and improves, requires new ways of measuring and circumventing it, and, therefore, remains a key focus for those studying censorship.

The concentrated research attention on countries like China may lead to concerns about "overstudying" specific regions. For instance, over the last decade, numerous studies have repeatedly analyzed well-documented censorship methods, such as DNS-based censorship in China [21,28,32,43,77,127,141,155,158,175]. While many of these studies have advanced the field by refining measurement methodologies and uncovering new censorship behaviors, others have simply reiterated previously established findings, such as the ongoing use of DNS manipulation. This repetition risks diverting attention and resources away from understudied regions.

Censorship in Europe, including within the EU, does exist, contrary to common beliefs. A 2021 study on censorship in Spain [14] uncovered network interference at the DNS, HTTP, and HTTPS levels, implemented by major ISPs in the country. The network-level interference was also evident during the period surrounding the 2017 Catalan independence referendum. Using open-source data provided by OONI [37], the researchers found that websites related to the referendum were blocked by ISPs in the week leading up to the vote. While censorship mechanisms in Spain (and other EU countries) have primarily targeted websites

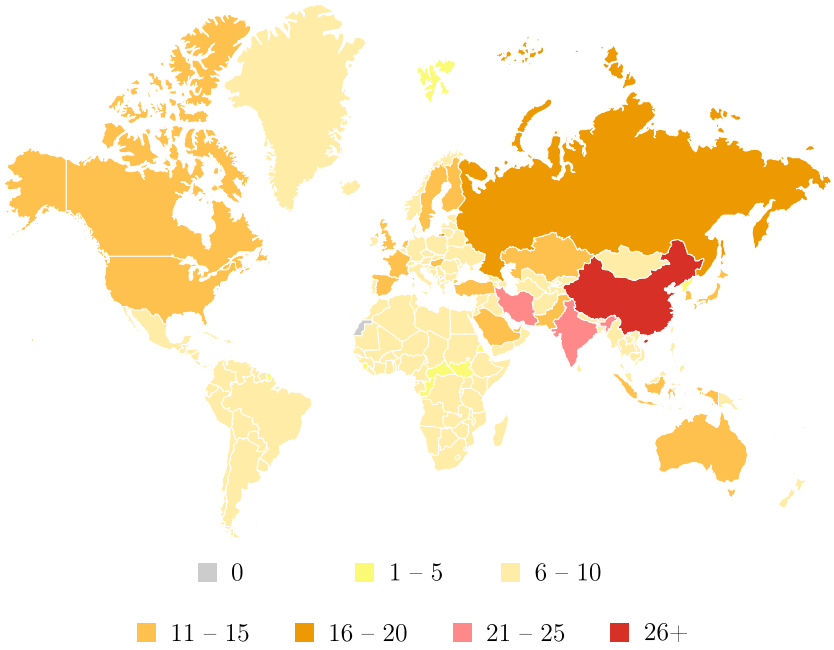

Fig. 3. Geographical focus of Internet censorship measurement studies (2015–2025).

related to gambling, copyright infringements, and other illegal content [189], the study demonstrated that these censorship mechanisms have also been repurposed to enforce other types of censorship, including political censorship.

6.2.2. Study duration

Internet censorship measurement studies vary significantly in duration. Table 2 presents the study durations in months for each study. Fig. 4 summarizes the study durations in a boxplot, based solely on studies that explicitly reported the length of their study period. The median study duration is 6 months (avg: 11.07 months), and 50% of the studies report measurement durations between 2 and 13.5 months. Among 50 studies, five studies stand out as outliers due to their long study durations: [115] analyzed over 86M OONI [37] data records related to outbound traffic throttling in Russia over a three-year period; [14] analyzed >3M OONI [37] data points focusing on websites advocating for civil rights in Spain over a period of 47 months; [33] identified correlations between Internet shutdowns and days of elections, protests, and coups using a 48-month dataset of ICMP, IP, and BGP data from 155 countries and territories; [32] examined over 70 billion Censored Planet [38] data points across 48 months; [34] conducted a 60-month-long study on 77 countries and territories to identify different blocking methods.

Since censorship practices are constantly evolving, the duration of a study can greatly influence its findings. Studies with very short measurement periods may capture only temporary or isolated instances of censorship. Although such studies still provide important contributions by highlighting these brief phenomena, their results may not generalize well to long-term trends, and they can quickly become outdated. This is especially relevant given that many studies either capture only a brief snapshot [3,23,69,127,129,144,155,174,181] or do not specify their exact measurement duration, such as [27,28,43,44,113,118,128,132,143,150] (indicated as "–" in Table 2).

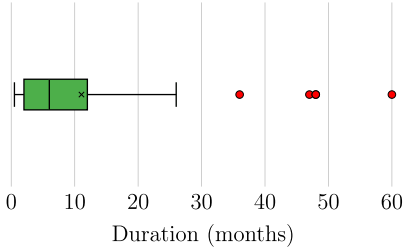

Fig. 4. Duration of Internet censorship measurement studies.

6.2.3. Network protocol prevalence

Researchers investigating Internet censorship have focused on different network protocols, and their prevalence in measurement studies offers insight into common censorship targets. Fig. 5 illustrates this by presenting the distribution of network protocols across 76 different Internet censorship measurement studies. The protocols are organized by their OSI layer, and the size of each bubble corresponds to the number of studies that specifically examined that protocol. Many studies, such as [13,18,27,32,33,42,43,124,128,137,144,150,151] focus on multiple protocols, some of which are interdependent.

This means that if a censor implements IP-based blocking, it will implicitly affect protocols that rely on IP, such as TCP or other higher-level protocols. For example, although most SSH traffic operates over TCP port 22 and can be censored by simply blocking that port, it is treated as a distinct protocol in Fig. 5 because researchers in [18,163] specifically focused on SSH in their investigations.

The most commonly used protocols in the context of Internet censorship are TLS, DNS, and HTTP. This focus is expected, as these are fundamental protocols upon which the majority of Internet services rely. Because these protocols are so widely used, censors can perform fine-grained filtering with minimal risk of causing collateral damage to unrelated services.

At the OSI application layer, 41 studies used TLS or HTTPS-based measurement probes, which are designed to detect TLS SNI-based filtering. This form of censorship relies on the TLS SNI field found in the TLS Client Hello message, which reveals the intended destination domain. By inspecting this field, censors can block specific services selectively. Similarly, 42 studies utilized DNS-based measurement techniques. These methods involve sending DNS queries to resolvers, where the queries containing the target domain names can be intercepted and filtered by censors. HTTP-based measurements were the most common, with

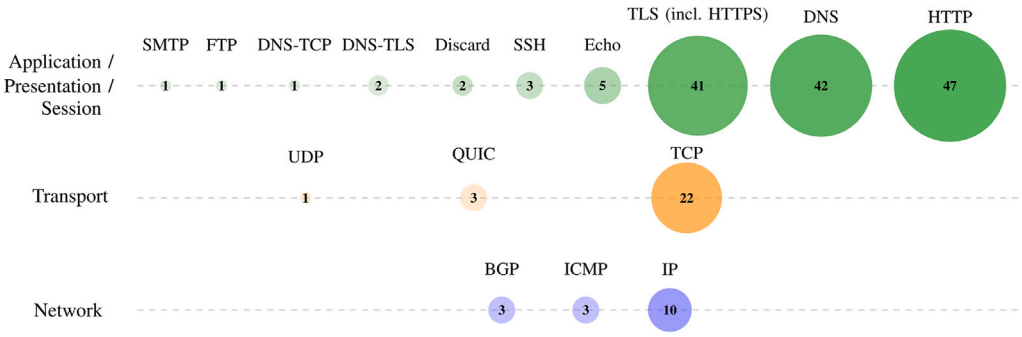

Fig. 5. Prevalence of network protocols used for Internet censorship measurements, arranged by size and OSI layer.

47 studies sending HTTP traffic to target servers to determine whether messages were being blocked or tampered with.

At the OSI transport layer, 22 studies investigated censorship based on TCP port blocking. One paper [116] investigated pure UDP traffic dropping by middleboxes in Iran. Additionally, three studies examined QUIC censorship by comparing the reachability of QUIC-capable domains over "traditional" TLS-over-TCP and over QUIC. These papers [111,150,154] aimed to identify cases where domains accessible via TLS-over-TCP were blocked when accessed using QUIC. The authors of [111] found differences between the QUIC blocklists and existing TLS, DNS, and HTTP blocklists in China's GFW and uncovered that QUIC *Initial* packet decryption degrades the performance of the firewall.

At the OSI network layer, 10 studies performed IP-level reachability measurements to detect outright IP blocking. Furthermore, three studies [15,33,176] focused on BGP disruptions induced by censors. For example, [176] analyzed topological data from 219 countries and territories and identified 108 censoring ASes, 32 of which impose censorship affecting users outside their own jurisdictions due to leakage of censorship policies through BGP.

6.2.4. Proportion of longitudinal studies

In general, longitudinal studies "*employ continuous or repeated measures*" [190] over long periods of time. In the context of Internet censorship, longitudinal studies aim to capture the "*volatility and spatiotemporal variability of Internet censorship*" [2]. This application of longitudinal methodology is particularly important for understanding how censorship strategies evolve across different times and geographic regions. Longitudinal Internet censorship studies, therefore, not only need to be of a certain duration, but also require repeated measurement of the same endpoints, using the same measurement techniques and the same vantage points, to test for spatiotemporal variability.

Fig. 6 presents the proportion of the total number of reviewed studies (red) as compared to the number of longitudinal studies (blue) over time.

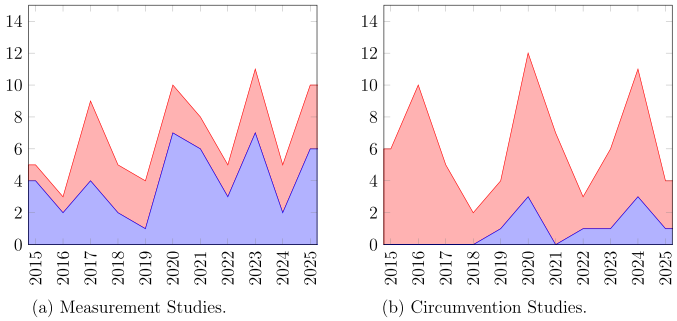

(a) Measurement Studies. (b) Circumvention Studies.

Fig. 6. Proportion of total number of reviewed studies (red) and number of longitudinal studies (blue) over time (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.).

Fig. 6(a) shows that the proportion of longitudinal measurement studies is high, whereas censorship circumvention studies commonly do not test the efficacy of their proposed methods over longer periods of time (see Fig. 6(b)). This can be partly attributed to the practical challenges associated with long-term testing, such as maintaining the circumvention infrastructure or adapting to new censorship tactics. However, more recently, some authors who proposed new censorship circumvention methods have also accompanied their studies with longitudinal evaluations and live deployments [99,191–195].

As mentioned in Section 7.2.3, [99] proposed Snowflake, a proxying system using temporary proxies that accept peer-to-peer WebRTC protocols to circumvent censorship. At the time of publication, Snowflake [98] had already been deployed in the Tor Browser and Orbot for multiple years. It has already proven to be a successful circumvention tool by helping thousands of users to bypass restrictions during network disruptions in Russia and Iran. By March 2024, Snowflake was supporting an estimated average of 35,000 concurrent users [99]. This longitudinal field-test is particularly valuable, because it provides real-world evidence of the tool's censorship circumvention capabilities. Similarly, in 2019 [45] presented Geneva (Genetic Evasion), a censorship evasion tool which automatically identifies evasion strategies. Since then, multiple studies, including [44,161,195,196], have extended the tool with new functionality. Over the years, Geneva was used to derive many successful censorship evasion strategies that were repeatedly tested against real censors.

Not all censorship measurement studies are longitudinal. There are exceptions, such as relatively short studies [69,144,146,155,174,181]. For example, [144] conducted a global study on connection tampering using Cloudflare CDN traffic. They collected and analyzed connection tampering signatures and found 19 unique fingerprints. As part of their study, the authors also investigated Iranian censorship over a 17-day period in September 2022, which had been sparked by protests. Since the connection tampering signatures were repeatedly collected on an hourly basis using the same vantage points and measurement methods, such studies, albeit short, can also be considered longitudinal.

6.3. Large-scale internet censorsHIP measurement platforms

Researchers and activists have been building large-scale Internet censorship measurement platforms [37–39,125,126] in recent years. These platforms capture and store censorship events from most countries and territories worldwide, and since they are open-source they provide a valuable resource for researchers and activists.

These platforms allow for censorship event correlation and deeper analysis of the underlying techniques. For instance, as of 2025, OONI [37] checks for DNS and TCP/IP tampering, whether instant messaging services (*e.g.,* WhatsApp, Signal, Telegram) are blocked, it further checks whether Tor (and Tor Snowflake [98]) is accessible. Moreover, OONI examines whether Psiphon [123] and RiseupVPN [197] are functional. In addition, OONI also provides speed and performance checks

and functions that help to detect the presence of middleboxes on a network.

Censored Planet [38] performs measurements using IP, TCP, DNS, HTTP, HTTPS, Echo and Discard [2]. At time of publication in 2020 [2], the authors had already collected and analyzed 21 billion data points from 95,000 vantage points and found censorship activities in over 100 different countries and territories. Notably, Censored Planet identified a 21-day period in 2019, during which Kazakhstan partially performed an HTTPS Interception attack on 37 unique domains [16]. The interception points were located using a TTL-based technique. In a different study [24], the platform was able to measure the characteristics and location of Twitter traffic throttlers using in-country vantage points and public data covering 401 unique Russian ASes. In another case study [70] on Russia, Censored Planet was used to find and compare differences in blocking behavior between data centers and residential ISPs. The authors acquired block lists from Russian activists and checked for TCP/IP blocking, keyword-based blocking and DNS manipulation. They identified higher levels of censorship and more block page injections in residential vantage points.

ICLab [39] is a platform designed to study Internet censorship using commercial VPNs located in different countries. Between 2016 and 2019, ICLab conducted censorship measurements in 62 countries using 81 vantage points and covering 234 different ASes [36]. ICLab can detect DNS tampering, IP-based blocking, and TCP packet injection. It also implements block page detection, which uses a set of regular expressions to detect block pages in HTTP responses. ICLab discovered 48 previously unknown block pages and also uncovered other forms of network interference, which are clearly distinguishable from normal network errors [36]. ICLab stores raw data to allow for future analysis and aims to minimize errors and the need for manual checks. This data also helped to reveal regional differences in censorship methods and showed how censorship in countries like India and Turkey changed alongside political events.

Two platforms specifically focusing on the “Great Firewall of China” (GFW) are GFWatch [77,125] and GFWeb [31,126]. GFWatch was introduced in 2021 and is a longitudinal censorship measurement platform which focuses on testing the DNS filtering behavior of China’s GFW. The platform operates a vantage point in the US, sends DNS requests to two controlled machines in China, and records the forged DNS responses from the GFW. The vantage points in China are used to verify that the censored domains are also censored inside China. At the time of publication, GFWatch had already probed 534 million distinct domains and identified 311 thousand censored domains. In addition, GFWatch identified a set of IPv4 and IPv6 addresses that are used to inject forged DNS responses. The authors also found a small number of websites that are unintentionally censored due to “greedy” regular expressions (regex) [77].

The evolution of GFWatch, called GFWeb [31], conducts large-scale measurements of the GFW’s HTTP and HTTPS filtering capabilities. GFWeb uses bidirectional, continuous measurements to assess the GFW, enabling cross-protocol comparisons with GFWatch’s DNS-based measurements. In total, GFWeb tested 1.02 billion domains, identifying censorship of approximately 943,000 domains and 55,000 Pay-Level Domains (PLD). A notable finding is that the GFW’s censorship is asymmetrical. That is, many domains are censored when accessed from within China, but the same filtering is not triggered when those domain names are included in requests sent from outside China through the GFW to destinations inside the country [31]. All datasets collected by GFWatch and GFWeb are publicly available [125,126].

7. Internet censorship circumvention research

Similarly to Section 6, this section classifies and synthesizes contemporary research on Internet censorship circumvention. First, it provides a chronological synthesis of the literature, which is based on 70 studies presented in Table 3. Then it examines the functional characteristics of circumvention tools, including the proportion that support bootstrapping versus those intended solely for operational use, as well as the degree to which tools are application-agnostic. Lastly, a taxonomy of censorship circumvention techniques is presented in Fig. 9. The two main categories identified are routing-based and obfuscation-based circumvention. Routing-based circumvention techniques can be further divided into Proxying, Alibi Routing, and Connection Splitting. Obfuscation-based solutions can be divided into Steganography, Protocol Mimicry, Deep Packet Inspection (DPI) Evasion, and Covert Tunneling. Each subcategory is discussed in a separate subsection.

7.1. Analysis

In total, 70 Internet censorship circumvention studies (see Table 3) were analyzed with respect to the following aspects:

- *Protocols* refers to the Internet protocols used by the circumvention traffic as it traverses censorship infrastructure (*i.e.,* the protocols observable by an on-path censor).
- *Type* denotes the classification of the circumvention method, following the taxonomy outlined in Fig. 9.
- *Throughput* indicates the highest throughput reported in at least one experiment conducted by the authors. If left blank, it means that neither throughput nor goodput was specified, or that throughput is uncapped and varies significantly depending on the type of transmitted data. It is important to note that the throughput of a circumvention system is inherently constrained by the bandwidth of the underlying communication channel. As a result, the throughput values have only limited expressiveness.
- *Longitudinal* refers to the type of study conducted. Longitudinal Internet censorship circumvention studies have repeatedly tested the effectiveness of their circumvention techniques over longer periods of time against multiple adversaries.
- *Bootstrapping* indicates whether the circumvention system addresses the bootstrapping problem—that is, how users obtain the initial information (*e.g.,* proxy addresses or shared secrets) required to connect to the service. Systems that rely solely on out-of-band channels (*e.g.,* email, encrypted messengers) for this purpose are marked accordingly. There are also systems that “only” provide solutions to the bootstrapping problem, without presenting a fully functional circumvention tool, or solutions which specifically optimize the bootstrapping process and do not focus on the operational phase [193,198], which have also been classified as “bootstrapping-only” solutions.
- *Operational* use means that, in contrast to bootstrapping, the circumvention tool also supports the ongoing transmission of data after the initial connection is established. This indicates it provides a stable and usable data channel with sufficient throughput to transmit information.
- *Application-Agnostic* censorship circumvention tools are designed to transmit arbitrary data and support a wide range of applications, independent of specific protocols or use cases. In contrast, non-agnostic tools are limited to particular applications or data types. A half circle in Table 3 means that the circumvention solution is only designed for a specific application but may transmit arbitrary data (*i.e.,* partly application-specific). For example, Telepath [199] is only designed for the video game *Minecraft* and Rook [200] for *Team Fortress 2*, a first person shooter that depends on privately-hosted servers.
- *Real-World Evaluation* refers to papers that have conducted extensive experiments involving countries that enforced Internet censorship at the time of publication. It means that researchers have tested their circumvention solutions against real middleboxes of censors. The gold standard in this category is a longitudinal study across multiple censorship regimes. This column focuses exclusively on real-world censorship evasion effectiveness and does not account for evaluations of other metrics such as resource usage or

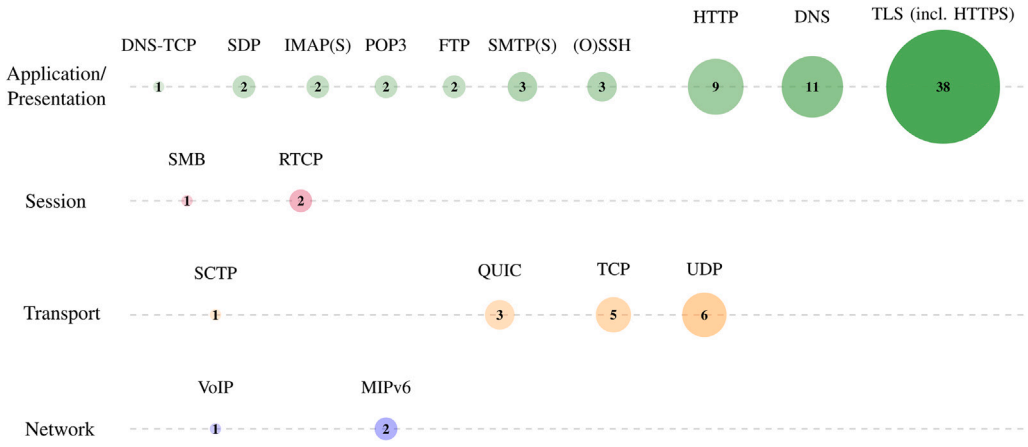

Fig. 7. Prevalence of network protocols used for Internet censorship circumvention, arranged by size and OSI layer.

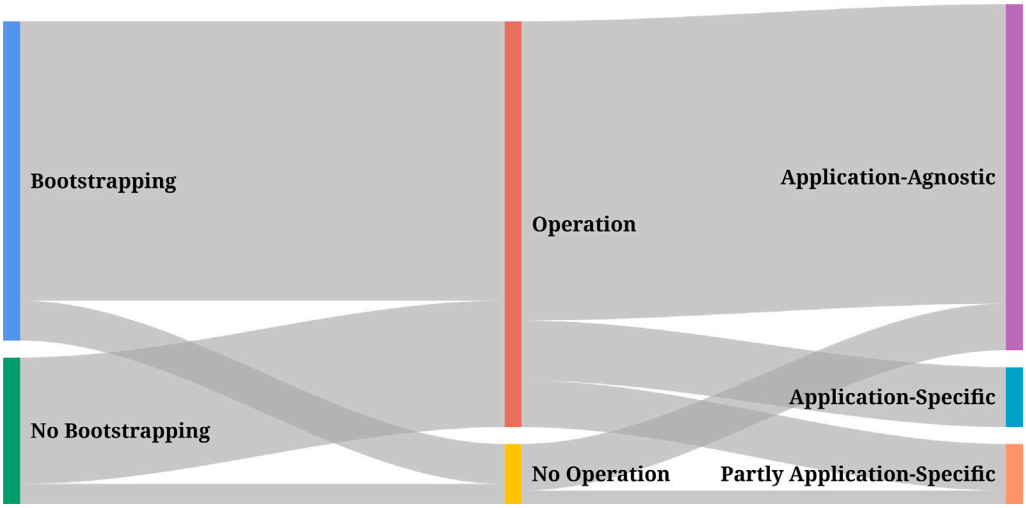

Fig. 8. Sankey diagram of circumvention tools by intended use and application specificity.

latency. While some studies [201,202] use simulations to approximate real-world conditions, these are not considered real-world evaluations under this criterion. There are also theoretical approaches (*e.g.,* [203]) that achieve provable undetectability, which may lessen the need for empirical testing.

- *Tool* points to the implementation of the circumvention technique. It has been marked as "N/A" if there was no link provided to an implementation, no GitHub repository, and no entries on other open-source platforms, such as *paperswithcode.com*.

7.1.1. Network protocol prevalence

Fig. 7 presents all Internet protocols that are used in the reviewed censorship circumvention literature. For each circumvention method, the Internet protocols visible from the perspective of the censor (*i.e.,* the on-path observers) are listed in Table 3. Some circumvention tools rely on multiple protocols; although only some protocols may be integral to the actual circumvention process, all involved protocols are listed. As in Fig. 5, protocols that depend on each other (for example, SSH depends on TCP) are displayed as separate protocols in Fig. 7. Protocols that could not be classified as network protocols or entire suites of protocols (*e.g.,* WebRTC) were not included in the bubble diagram.

Similarly to the protocols used for measuring censorship, TLS is also used in the majority of censorship circumvention tools. Some categories, such as Refraction Networking rely exclusively on TLS because the use of TLS is inherent to the method's operation [236,246,257,261,266,270–272,277,279,288,298]. Others are protocol-agnostic, such as provably-secure steganography [203], or offer a way to dynamically choose a protocol, *e.g.,* [82,217,222,255].

Overall, there is a clear overlap between the OSI application-layer protocols used to circumvent censorship and those used to measure it. This overlap is logical: the protocols most commonly used to access information are the ones targeted by censors. As a result, these same protocols become the focus of researchers who measure censorship.

7.1.2. Tool-specific functionalities

Fig. 8 presents a Sankey diagram showing *(i)* the proportion of circumvention tools that do or do not incorporate a bootstrapping mechanism, *(ii)* the proportion of tools that were designed for operational use and those designed exclusively for bootstrapping, and *(iii)* the proportion of tools that are application-agnostic, application-specific, or partly application-specific.

The diagram shows that the majority of circumvention tools that implement a bootstrapping mechanism were also designed for operational use, *i.e.,* they provide enough throughput and reliability to establish a usable communication channel, rather than just focusing on initial connection setup. Around 30% of the tools do not implement a dedicated

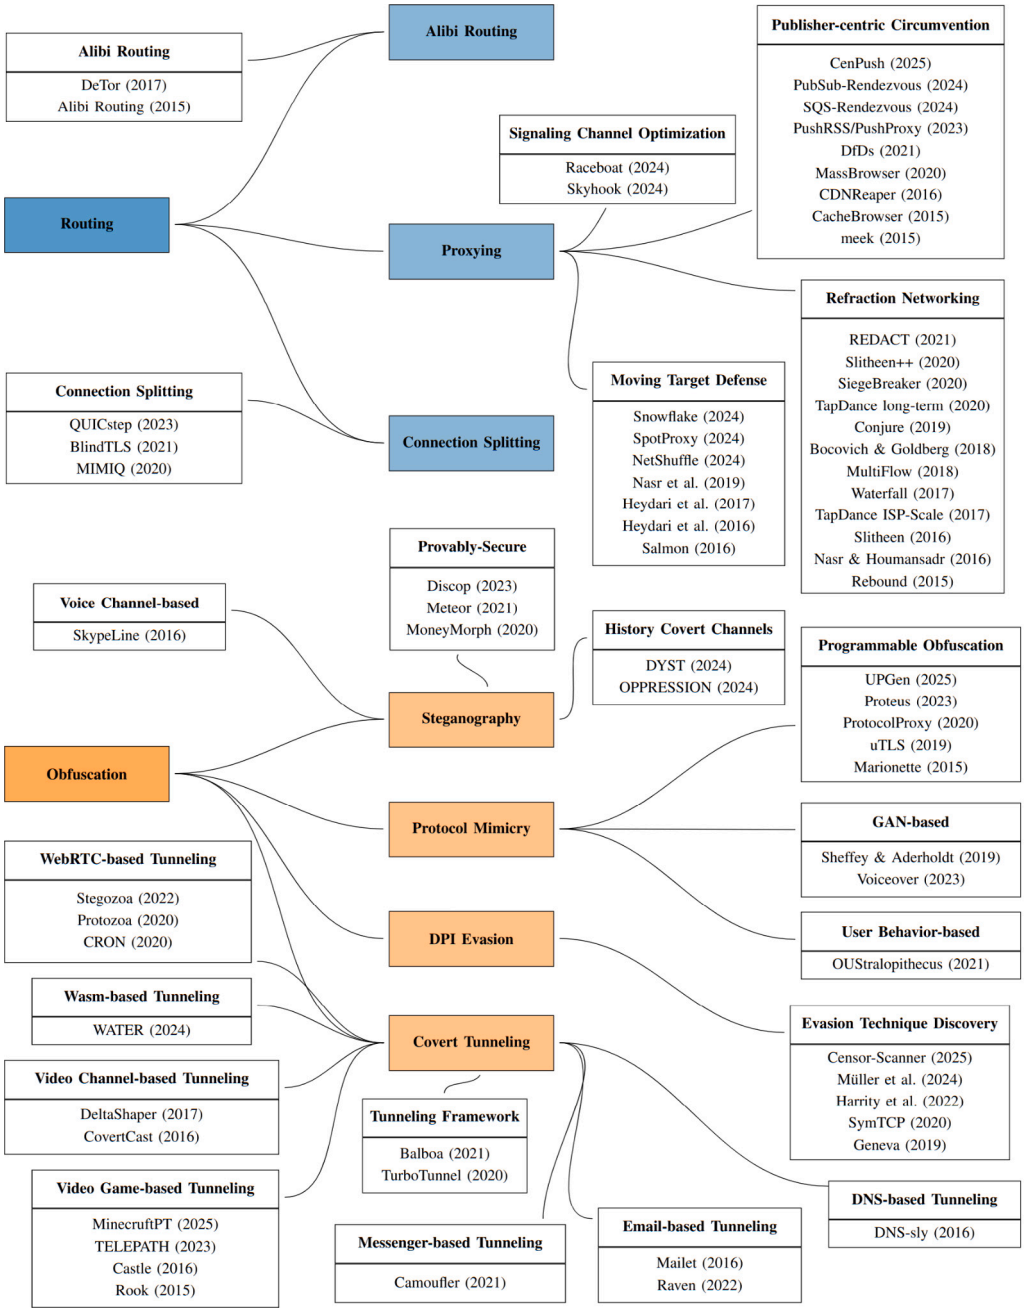

Fig. 9. Taxonomy of Internet censorship circumvention techniques (2015–2025).

bootstrapping mechanism, and rely on out-of-band bootstrapping. Only two tools were designed neither for operational nor bootstrapping use. One of them, Alibi Routing [201], is a peer-to-peer overlay routing system that does not provide circumvention functionality per se, but rather proofs of avoidance of certain geographic regions.

Most tools that were designed for operational use are also application-agnostic, meaning they can transmit arbitrary data and are not dependent on a specific application. A small number of tools are application-specific. For example, publisher-centric circumvention systems that use CDNs to buffer specific types of content (*e.g.,* [284,297]). Using such systems, users can only circumvent censorship to retrieve pre-selected content. Other application-specific circumvention tools include systems that depend on a third-party program, such as [294], which depends on Skype.

7.2. Proxying

This subsection surveys a range of approaches which are based on Refraction Networking, Publisher-Centric Circumvention, Moving Target Defense, and Signaling Channel Optimization.

7.2.1. Refraction networking

Refraction Networking is a type of proxying technique that employs ISP-level routers to facilitate censorship circumvention. It involves installing the circumvention software directly on the Internet routers of volunteer ASes, known as decoy ASes. Refraction Networking makes use of end-to-middle proxying between an AS and the target server [272].

Refraction Networking was introduced in 2011, under the name of Decoy Routing. Telex [302], Cirripede [83], and Curveball [303] were among the first proposed systems that implemented this circumvention

Table 3
Chronological synthesis of relevant internet censorship circumvention research (2015-2025).

Work	Year	Protocols	Type	Throughput	Longitudinal	Bootstrapping	Operational	Application-Agnostic	Real-World Evaluation	Tool
[191]	2025	TLS, Tor	Publisher-Centric Circumvention	–	●	○	●	○	●	CenPush [204]
[205]	2025	flexible	Protocol Mimicry	–	○	●	●	●	●	UPGen [206]
[207]	2025	TLS	DPI Evasion	–	○	●	●	●	●	Censor-Scanner [208]
[209]	2025	TCP	Covert Tunneling	350 Kbps	○	○	●	◐	◐	MinecruftPT [210]
[211]	2024	flexible	Steganography / Covert Tunneling	–	○	○	●	●	○	DYST [212]
[82]	2024	*e.g.,* HTTPS, Snowflake, obfs; (flexible)	Proxying (Framework)	–	○	●	○	●	○	Raceboat [213]
[214]	2024	HTTPS	Proxying / Covert Tunneling	–	○	●	○	◐	○	Skyhook [213]
[215]	2024	HTTP	DPI Evasion	–	○	●	●	●	●	[216]
[217]	2024	flexible	Steganography / Covert Tunneling	–	○	○	●	●	○	OPPRESSION [218]
[192]	2024	HTTPS, DNS	Shuffle Proxying	–	●	○	●	●	●	NetShuffle [219]
[99]	2024	SDP, WebRTC, TurboTunnel, TLS (TOR)	Proxying / Moving Target Defense	–	●	●	●	●	●	Snowflake [98]
[220]	2024	Wireguard, Snowflake	Proxying / Moving Target Defense	~ 33 Mbps	○	○	●	●	○	SpotProxy [221]
[222]	2024	*e.g.,* TLS, Turbotunnel; (flexible)	Wasm-based Circumvention Deployment	1830 Mbps	○	○	○	●	○	WATER [223]
[198]	2024	HTTP, HTTPS (Google Pub/Sub)	Publisher-Centric Circumvention	–	○	●	○	●	○	PubSub-Rendezvous [224]
[193]	2024	SDP, WebRTC, TurboTunnel, TLS (TOR)	Publisher-Centric Circumvention	–	●	●	○	●	●	SQS-Rendezvous [225]
[226]	2023	QUIC	Covert Tunneling	–	○	●	●	●	○	QUICstep [227]
[199]	2023	TLS	Steganography	1300 Kbps	○	●	●	◐	○	TELEPATH N/A
[194]	2023	HTTP, HTTPS, UDP	Publisher-Centric Circumvention	4.86 Mbps	●	●	●	●	●	PushRSS, PushProxy N/A
[228]	2023	TCP (flexible)	Protocol Mimicry / Covert Tunneling	–	○	○	●	●	○	Proteus [229]
[230]	2023	UDP, RTCP (*e.g.,* Skype)	Protocol Mimicry / Covert Tunneling	62.32 Bps	○	●	●	◐	○	Voiceover [231]
[203]	2023	Protocol-agnostic	Steganography	–	○	○	●	●	○	Discop [232]
[195]	2022	HTTP, DNS	DPI Evasion	–	●	●	●	●	●	Geneva [107]
[233]	2022	SMTPS, IMAPS	Covert Tunneling	–	○	●	●	●	◐	Raven N/A
[234]	2022	WebRTC	Covert Tunneling / Steganography	8.2 - 23.8 Kbps	○	●	●	●	○	Stegozoa [235]
[236]	2021	TLS	Refraction Networking	–	○	●	●	●	○	REDACT N/A
[237]	2021	HTTPS	Protocol Mimicry / Covert Tunneling	2.2 Mbps	○	●	●	●	○	OUStralopithecus [238]
[239]	2021	flexible	Steganography	–	○	○	●	●	○	Meteor [240]
[241]	2021	TLS	Proxying / Covert Tunneling	–	○	●	●	●	●	BlindTLS N/A
[242]	2021	TLS	Covert Tunneling / Traffic Obfuscation	–	○	●	●	●	○	Balboa [243]
[244]	2021	Signal, TLS, MTProto	Covert Tunneling	–	○	●	●	○	◐	Camoufler N/A
[245]	2021	DNS, HTTPS	Publisher-Centric Circumvention	–	○	●	●	○	○	DfDs N/A
[246]	2020	TLS	Refraction Networking	~ 600 Kbps	○	●	●	●	○	Slitheen++ [247]
[248]	2020	TCP	Steganography	–	○	●	○	◐	○	MoneyMorph [249]

(*continued on next page*)

Table 3 *(continued)*

Work	Year	Protocols	Type	Throughput	Longitudinal	Bootstrapping	Operational	Application-Agnostic	Real-World Evaluation	Tool
[250]	2020	QUIC	Covert Tunneling	–	○	●	●	●	○	MIMIQ [251]
[252]	2020	WebRTC	Covert Tunneling	–	○	●	●	●	○	CRON N/A
[196]	2020	DNS-over-TCP, FTP, HTTP, HTTPS, SMTP	DPI Evasion / Covert Tunneling	–	●	●	●	●	●	Geneva [107]
[253]	2020	WebRTC	Covert Tunneling / Steganography	1.4 Mbps	○	●	●	●	●	Protozoa [254]
[255]	2020	*e.g.,* KCP, QUIC, SCTP & obfs4, meek; (flexible)	Covert Tunneling	–	○	○	●	●	○	TurboTunnel [256]
[257]	2020	TLS	Refraction Networking	–	○	○	●	●	○	SiegeBreaker [258]
[259]	2020	TCP, UDP	Proxying / Publisher-Centric Circumvention	10–30 Mbps	●	●	●	●	○	MassBrowser [260]
[261]	2020	TLS	Refraction Networking	100 Kbps	●	●	●	●	●	TapDance (long-term) [262]
[263]	2020	TCP	DPI Evasion / Covert Tunneling	–	○	●	●	●	●	SymTCP [264]
[265]	2020	UDP	Protocol Mimicry / Covert Tunneling	182 Bps	○	●	●	●	○	Protocol Proxy N/A
[101]	2019	HTTPS	Protocol Mimicry / Traffic Obfuscation	–	○	○	●	●	○	[102]
[266]	2019	TLS, OSSH	Refraction Networking	–	○	●	●	●	●	Conjure [267]
[85]	2019	TLS (TOR)	Proxying / Moving Target Defense	–	○	●	○	●	○	N/A
[268]	2019	TLS	Protocol Mimicry / Traffic Obfuscation	–	●	●	●	●	●	uTLS [269]
[270]	2018	TLS	Refraction Networking	–	○	●	●	●	○	N/A
[271]	2018	TLS	Refraction Networking	–	○	●	●	●	○	MultiFlow N/A
[272]	2017	TLS	Refraction Networking	–	○	○	●	●	●	Waterfall [273]
[274]	2017	UDP, RTCP (*e.g.,* Skype)	Covert Tunneling / Steganography	7.2 Kbps	○	●	●	◐	○	DeltaShaper [275]
[276]	2017	MIPv6	Proxying / Moving Target Defense	–	○	●	●	●	○	N/A
[277]	2017	TLS	Refraction Networking	–	○	●	●	●	●	TapDance (ISP-Scale) [262]
[202]	2017	TLS (TOR)	Alibi Routing	–	○	●	●	●	○	DeTor [278]
[279]	2016	TLS	Refraction Networking	–	○	●	●	●	○	Slitheen [280]
[281]	2016	HTTPS	Covert Tunneling / Steganography	169 Kbps	○	○	●	○	○	CovertCast [282]
[283]	2016	MIPv6	Moving Target Defense	–	○	●	●	●	○	N/A
[284]	2016	DNS, HTTP, HTTPS	Publisher-Centric Circumvention	–	○	●	●	○	●	CDNReaper [285]
[286]	2016	SMTP, IMAP, POP3	Covert Tunneling / Proxying	–	○	●	●	◐	○	Mailet [287]
[288]	2016	TLS	Refraction Networking	–	○	○	●	○	○	N/A
[289]	2016	DNS	Covert Tunneling	–	○	○	●	●	●	DNS-sly N/A
[290]	2016	UDP	Covert Tunneling	320 Bps	○	○	●	◐	○	Castle [291]
[292]	2016	DNS, HTTP, HTTPS, SSH	Proxying / Moving Target Defense	1.5 Mbps	○	○	●	●	○	Salmon [293]
[294]	2016	VoIP	Steganography	2400 Bps	○	●	●	○	○	SkypeLine N/A
[200]	2015	OTR	Covert Tunneling / Steganography	30 Bps	○	○	●	◐	○	Rook [295]
[201]	2015	–	Alibi Routing	–	○	○	○	●	○	Alibi Routing [296]
[297]	2015	DNS, HTTP, HTTPS	Publisher-Centric Circumvention	–	○	●	●	○	●	CacheBrowser [285]
[86]	2015	DNS, HTTPS	Publisher-Centric Circumvention	–	○	●	●	○	○	meek [100]
[298]	2015	TLS	Refraction Networking	126 Kbps	○	○	●	●	○	Rebound [299]
[300]	2015	HTTP, SSH, SMB, FTP, POP3	Protocol Mimicry (framework)	68.2 Mbps	○	●	●	●	○	Marionette [301]

technique. The authors of [272] developed a downstream-only decoy routing solution that offers resistance against rerouting attacks. This is due to the fact that censors have less control over downstream BGP routes [272]. It is initiated by opening a TLS connection from a client within a censored region to an overt destination via a decoy ISP router (*i.e.,* refraction station), typically involving steganographic tags which are used to register clients with the refraction station. The refraction station observes the signal, parses the tag, and opens both a decoy connection to an overt destination and a connection to the censored destination. Telex [302], for instance, hides the steganographic tag in the TLS Client Hello message sent to the overt destination. Cirripede [83] embeds the tag in the Initial Sequence Numbers (ISN) of TCP SYN messages. TapDance [277] and Conjure [266] encode the registration signal, so that it is indistinguishable from random bytes, and embed the data in HTTPS request body to hide the signal from the censor.

Unlike traditional proxying, which relies on contacting servers that perform the forwarding, refraction networking is an "en route" proxy. A key advantage of this approach is that censors cannot easily block refraction networking using conventional IP-based censorship, as there are no fixed proxy endpoints to target. As a result, a major challenge for censors attempting to counteract refraction networking is the high risk of collateral damage. Blocking these routers would disrupt not only circumvention efforts but also a significant portion of regular Internet traffic. Hence, censors are de-incentivized from outright banning decoy routers, as doing so would negatively impact their own networks and users.

From an adversarial perspective, there is also a technique which attempts to bypass decoy routing by strategically altering the network path – also called a Routing Around Decoys (RAD) attack [304]. In this attack, censors manipulate BGP routing to ensure that user traffic never traverses ISP-level routers hosting decoy nodes. By rerouting connections away from participating ISPs, censors can effectively prevent users from leveraging refraction networking. Countermeasures against RAD attacks include increasing the number of participating ISPs and deploying decoy routers in more diverse network locations.

However, increasing the number of participating ISPs is not practical, as shown in [305]. The authors provided an analysis of the placement of decoy routers in the Internet. They demonstrated that intercepting the majority of network paths to popular global sites from censored regions can be achieved by deploying decoy routers in approximately 30 key ASes. They identified that, despite the small number of ASes needed, an effective refraction networking system would still require about 11,700 decoy routers, making deployment a significant logistical effort. The substantial financial cost of such a system also limits its practical feasibility.

7.2.2. Publisher-centric circumvention

Content Delivery Networks (CDNs) provide cached content at the network edge. CDNs are strategically placed at Internet eXchange Points (IXPs) to reduce latency for users who access web content. CDNs also use other optimizations, such as load balancing, TLS connection reuse, and various compression techniques to improve performance and reduce bandwidth consumption [306].

Researchers and activists have found CDNs useful in circumventing censorship (see Table 3) because they distribute content across multiple, geographically dispersed servers and host different types of content behind a shared domain name. This design makes it difficult for censors to block specific websites without also affecting other services. Since blocking a major CDN could lead to widespread Internet disruptions, governments often hesitate to take such drastic measures; hence, they indirectly allow censored content to remain accessible through CDNs.

One example of a CDN-based circumvention method is *Domain Fronting*, which was first introduced by [86] in 2015. To route a request to a destination server it relies on the fact that CDNs use the HTTP Host header, which is encrypted and therefore invisible to a middlebox, instead of the TLS Server Name Indication (SNI), which is visible in cleartext to a middlebox.[1] The DNS request preceding the TLS handshake also carries the "front domain"; therefore it is impossible for the censor to block the circumvention traffic without also blocking the "benign" front domain [86].

Domain fronting has already been restricted by Big Tech companies like Microsoft and Google, reducing its availability as an anti-censorship tool [87,88]. Due to concerns over misuse and potential policy violations, other companies, such as Cloudflare, also refrain from offering domain fronting as a feature [88].

Domain Shadowing, a technique introduced by [245], also leverages the fact that users can add rules in their CDN account configuration to map a specific (uncensored) HTTP Host header to a (censored) target Host header. Domain shadowing is also a term used for a method in which malicious actors compromise a domain administrator's account to create multiple subdomains. This allows them to host malicious pages that evade detection and bypass blacklists [307]. However, in the context of Internet censorship research, the term domain shadowing has taken on a different meaning: Domain Shadowing, as compared to Domain Fronting, does not require the user to modify the TLS SNI value nor the HTTP Host header, but rather requires rewriting rules in a personal CDN account. The rules in the CDN account rewrite the HTTP Host header of a request and forward it to a censored destination. From the perspective of a censor, domain shadowing appears as normal, non-suspicious traffic because all visible indicators suggest that the user is accessing an allowed domain. Specifically, the censor sees the DNS request for the shadow domain, the TLS SNI, and the destination IP address pointing to the CDN.

More broadly, publisher-centric circumvention also includes all approaches that leverage publicly available Publish–Subscribe (Pub/Sub) messaging services. For instance, [198] developed a censorship-resistant rendezvous channel based on a Google Pub/Sub service. This channel helps users within censored regions connect to Tor bridges by relaying necessary information about bridges via the Pub/Sub service. Similarly, [193] used Amazon's Simple Queue Service (SQS) to establish a rendezvous channel, which has been integrated into Snowflake [98] and the Tor Browser.

Xue and Ensafi [194] developed a circumvention technique that uses Really Simple Syndication (RSS) aggregators to subscribe to content updates published on uncensored platforms. These updates are delivered as push notifications to a censored device and can serve as a channel for transmitting information necessary for the bootstrapping phase of another circumvention method.

Another publisher-centric circumvention technique is directly accessing cached content served by CDNs, which allows censored users to retrieve information without directly contacting the origin server. From a censor's perspective, the barrier to block entire CDNs is high, because doing so could potentially cause significant collateral damage.

Several CDN-based circumvention systems have been proposed, including CacheBrowser, which was designed by [297] and evades DNS-based censorship by fetching otherwise blocked content directly from CDN servers. CacheBrowser has been tested against China's GFW and was successful in circumventing censorship for forbidden websites, such as *facebook.com*.

CDNReaper, proposed by [284], is a CDNBrowsing system that mitigates the risk of CDNBrowsing fingerprinting attacks and is available as a Google Chrome and Firefox browser plugin. MassBrowser was introduced by [259] and is a proxy-based circumvention system run by a network of volunteers. It relies on *(i)* users with unrestricted access to the

[1] The SNI is only visible in TLS connections that do not use Encrypted Client Hello (ECH).

Internet to proxy blocked traffic for censored users and on *(ii)* censored users to forward traffic to others facing different censorship restrictions.

7.2.3. Moving target defense

Proxy enumeration attacks pose a challenge to circumvention systems. Censors can impersonate benign clients and discover circumvention proxies, which they can then block based on IP addresses [85]. To counteract proxy enumeration and blocking attacks, several proxy distribution and Moving Target Defense (MTD) systems have been proposed [85,99,192,220,276,283,292].

Douglas et al. [292] present a volunteer-based proxy distribution system named "Salmon", which classifies user behavior as trustworthy or suspicious based on their actions. As benign users participate in the network of proxies, they can improve their discrete trust levels using a referral-based system, while malicious users will be flagged and banned from participating in the network.

Heydari et al. [276,283] introduce a censorship-resistant method using MTD with Mobile IPv6 (MIPv6). Their system assigns users to access groups, where each access group is assigned to a pool of multiple MIPv6 care-of-addresses (CoAs). In MIPv6, CoAs are temporary IP addresses assigned to a mobile node when it is away from its home network. Assigning such temporary CoAs enables treating web servers as mobile nodes and hence allows them to rotate CoAs to evade blocking.

A game-theoretic approach to censorship-resilient proxy distribution was introduced by [85]. The authors modeled the proxy assignment problem as the *college admissions game* [308], in which students (*i.e.,* circumvention clients) compete against other students to be admitted to colleges (*i.e.,* proxies). The goal of the system is to help clients and proxies rank each other for optimal proxy assignment. The development of this game-theoretic approach was motivated by the fact that a significant proportion of existing proxies (*e.g.,* Tor bridges) are already blocked by censors, and new proxies can also be promptly blocked. Therefore, as suggested by [85], proxy "birth" intervals, central management of proxies, and the proportion between proxies and censoring agents play an important role in enhancing the censorship resilience of the proxy distribution system.

Temporary proxies have become popular circumvention approaches recently due to their ability to enhance blocking resistance and reduce the cost and complexity of proxy deployment. By implementing techniques such as reshuffling, rejuvenation, and rapid proxy rotation, these systems make it significantly harder for censors to enumerate and block proxies. Temporary proxies typically avoid reliance on fixed infrastructure. Instead, they use ephemeral nodes that frequently change their network identity—whether through dynamic IP reassignment, as seen in NetShuffle [192], or through temporary spawning of new proxy instances, as in SpotProxy [220] and Snowflake [99].

Another approach to increasing censorship resilience is to decouple physical proxy servers from the logical addresses they use. NetShuffle [192] is a modular service that implements this strategy by dynamically "shuffling" proxy servers, making it much harder for censors to block them based on IP addresses. This shuffling mechanism separates proxy nodes from fixed IP addresses by assigning them subdomain names within a domain owned by the edge network and dynamically mapping these subdomains to different IP addresses.

Snowflake, presented in [99], is a proxying system using temporary proxies that accept peer-to-peer WebRTC protocols. At the time of its publication, [99] presented the results of a three-year study in Russia, Iran, China, and Turkmenistan, during which it allowed many users to circumvent censorship deployed by these countries. Snowflake has been embedded into the Tor Browser [98]. The system uses a large, dynamic pool of lightweight proxies that can run in a web browser or as a browser extension.

SpotProxy, developed by [220], optimizes the cost of circumvention proxies by hosting low-resource and low-cost proxy instances in a public cloud. Similarly to NetShuffle, SpotProxy tries to minimize the time that one proxy uses a fixed IP address by constantly re-spawning new proxies with different IP addresses. To migrate between proxies without performance degradation, WireGuard and Snowflake [98] are used.

7.2.4. Signaling channel optimization

Vines et al. presented Raceboat [82] and Skyhook [214]. The former is a framework designed for managing multiple signaling channels and providing a mixing-and-matching feature to connect different unidirectional channels with one another. The latter is a cloud storage-based system, which allows for establishing short-lived signaling channels using Amazon AWS S3 buckets.

7.3. Alibi routing

Alibi Routing is a peer-to-peer overlay routing system introduced by [201], which is designed to provide users with proofs that their Internet traffic did not pass through specific predefined geographic regions—so called proofs of avoidance. Such proofs of avoidance are achieved through relay nodes positioned far enough from the forbidden area such that if the traffic had passed through both the relay and the region, the resulting delay would be detectable.

Alibi Routing requires no changes to existing Internet infrastructure, because it does not need the Public Key Infrastructure (PKI), any modifications to network protocols, or knowledge of the global routing topology. Instead, it relies on local measurements (like RTTs) and GPS-based location data. Tests on PlanetLab [309] and simulations show that the system works well for many source-destination pairs, especially when those points are far from the avoided regions [201].

However, the system also has some limitations. If the sender or receiver is very close to or inside the forbidden area, finding viable alibis becomes much harder or impossible. Alibi Routing is thus meant to complement, not replace, traditional censorship circumvention tools.

A proof of avoidance for the Tor overlay network is presented in [202]. The proposed technique, named "DeTor", allows a user to define specific geographic regions that should be excluded when creating a Tor circuit. Existing Tor options let users attempt to exclude certain countries, but in practice, this provides only the illusion of control because most circuits still route through those regions [202].

Instead of relying on methods like *traceroute* or network topology maps (which can be spoofed), DeTor uses geographic distance and speed-of-light constraints to prove where traffic could not have gone. DeTor includes two possible guarantees: *(i)* "Never-once": proves that packets never crossed a specified region and *(ii)* "Never-twice": ensures that packets do not pass through the same region at two different points in the route, which is important for stopping deanonymization attacks. Similarly to [201], DeTor also does not require any changes to Tor or the Internet's infrastructure and works using client-side RTT measurements. The authors showed that DeTor can successfully avoid certain regions like the U.S. in many cases [202].

7.4. Connection splitting

Connection Splitting is a routing-based censorship circumvention method which prevents censors from analyzing and blocking traffic. This method can either be implemented using custom circumvention protocols, or by leveraging existing multi-path functionality in protocols like QUIC or TLS (via session resumption).

A QUIC-based method, which uses frequent connection migrations to prevent an on-path observer from performing traffic analysis, is presented in [250]. MIMIQ splits a QUIC connection into smaller flows and chooses to migrate between these flows randomly, making it more difficult for censors.

BlindTLS, a circumvention method that leverages the TLS session resumption protocol, was proposed by [241]. The method first performs the TLS 1.2 handshake over an encrypted channel, which effectively hides the SNI value, and then resumes the connection over a standard

channel, using TLS session resumptions. The reason this method is useful is that either censors outright block TLS 1.3 Client Hello messages containing an ESNI, or ISPs drop TLS 1.3 handshakes since it is still an optional feature [241]. Hence, BlindTLS can help to evade TLS SNI-based censorship without compromising performance.

Similarly to [250], Jia et al. [226] introduce a method called QUICstep, which also relies on QUIC connection migrations. QUICstep first establishes a WireGuard VPN tunnel with the destination and sends the QUIC handshake over the WireGuard tunnel, which hides the TLS SNI value from a potential censor. Then a QUIC connection migration is triggered on the client-side, migrating the QUIC connection to a "normal" non-WireGuard network interface. The connection is then resumed and does not transmit any more handshake packets (*i.e.,* cleartext information).

7.5. Steganography

Steganography refers to the practice of hiding information within seemingly innocuous data to evade detection [310]. In the context of censorship circumvention, steganographic methods enable covert communication by embedding secret messages in ordinary traffic in ways that are resistant to surveillance or filtering.

This subsection surveys recent developments in steganographic techniques designed specifically for censorship resistance. The focus is on three distinct approaches: provably-secure steganographic schemes, which offer formal guarantees of undetectability; a "history covert channel" technique, which uses publicly available data as a medium for communication without modifying content; and a voice channel-based method, which embeds information within VoIP calls.

7.5.1. Provably-secure steganography

Provably-secure steganography refers to methods of hiding messages in cover media (like text, images, or other data) such that hidden communication is (provably) undetectable to adversaries. One example of a provably-secure steganographic bootstrapping scheme is MoneyMorph [248]. This scheme uses micro-transactions in blockchain-based cryptocurrencies to transmit the bootstrapping information, such as the IP addresses of proxy servers or cryptographic key material. The authors implemented the scheme on top of the Bitcoin, Zcash, Monero, and Ethereum networks and showed that the system is both fast and cheap: an encoding/decoding operation can be done in less than 50ms, costing less than 0.01 USD per transaction using Zcash [248].

Kaptchuk et al. [239] developed a censorship-resistant steganographic scheme that samples from generative models to mimic real human communication. The key challenge that [239] overcomes is the creation of stegotext that is indistinguishable from natural language communication, especially when the underlying communication channels (like English text) are complex and have highly variable entropy. Earlier methods either require impractically high entropy or lack provable security.

Similarly to mimicking the underlying distribution of natural language communication, [203] proposed Discop, a provably secure steganography system that samples from distribution copies of the original probability distribution of a message. Since the security of the system lies in the similarity of the probability distributions used, they employ the average and maximum values of the KL Divergence to quantify security. The method can embed up to 5.76 bits per token. Drawbacks of Discop are that the sender and receiver have to agree upon a shared steganography secret beforehand, such as a Pseudo Random Number Generator (PRNG) or a seed. This can be considered part of the bootstrapping process, albeit the authors do not explicitly describe the key exchange protocol. Arguably, Discop does not require real-world evaluation because it is a provably secure method that cannot be detected by steganalyzers, assuming that no secret key material is leaked.

7.5.2. History covert channels

Another solution to steganography-based censorship circumvention is presented by [211,217]. The authors of [217] introduce a new class of steganographic covert channels, called "history covert channels", which transmit secret information not by modifying or embedding data, but by pointing to existing, publicly accessible content on the Internet that matches the intended message (*i.e.,* a secret message). The method relies on an existing censorship-resistant channel for participants to exchange the address of the public website or data source to generate identical local codebooks. The participants then exchange small pointers referencing strings in this codebook, which represent the actual hidden messages to be communicated covertly. The evaluation methodology involved building a shared external codebook using sources like Twitter, Wikipedia, and digital books, then encoding secret messages as compact pointers to those sources. The findings show that system significantly reduces the amount of data transmitted and outperformed plain and gzip-compressed messages. When integrated with the Stegozoa [234] covert channel, it also improved throughput and reduced covert transfer times [217].

7.5.3. Voice channel-based

One example of a voice channel-based steganographic circumvention method is SkypeLine [294], which relies on Direct Sequence Spread Spectrum (DSSS) steganography and embeds secret information in VoIP traffic. The DSSS embedded information is only recoverable by parties possessing a pre-shared secret. The throughput of SkypeLine reaches up to 2400 bps using two newly designed modulation schemes (one binary, one m-ary). Unlike previous approaches, SkypeLine is compatible with any VoIP client and stays hidden both to the ear and in statistical analysis of the signal. It is also resistant to network interference due to an optional acknowledgment system that handles packet loss.

7.6. Protocol mimicry

Protocol Mimicry circumvention techniques aim to perfectly mimic ordinary network protocol traffic by, for example, replicating the statistical properties and protocol semantics of network packets, such as inter-arrival times, payload entropy, payload sizes and cleartext header information.

This subsection first briefly surveys programmable obfuscation techniques, which are very promising mimicry methods since they provide a flexible way of adapting to censorship strategies. Subsequently, this section provides examples of GAN-based and user behavior-based mimicry methods.

7.6.1. Programmable obfuscation

Traditional protocol mimicry selects a protocol and a certain version or implementation of it, and tries to mimic its statistical properties and protocol semantics as closely as possible in order to trick a censor into thinking that the traffic is legitimate communication using that protocol. In recent years, several programmable mimicry systems have been proposed [205,228,265,268,300], which automate the obfuscation process and dynamically adapt to the underlying network conditions and censorship strategies.

Marionette [300] is a network-traffic obfuscation framework designed to enable users to easily create and test a wide range of obfuscation strategies without modifying code or underlying systems. It allows users to modify the format of application-layer messages and adjust the statistical properties of connections. Marionette uses new template grammar and plugin systems to format messages so that they can evade DPI and pass through proxies and firewalls without detection.

Currently, TLS is likely one of the most prevalent protocols used for network traffic fingerprinting. Frolov & Wustrow [268] present uTLS, a library that, contrary to Marionette, focuses exclusively on mimicking

TLS fingerprints (*e.g.,* the order of the TLS cipher suites in the TLS client hello message). Since there is a range of different TLS implementations available, many of which significantly differ in terms of fingerprints, a specific TLS implementation may give away valuable information to an on-path observer. The authors analyzed over 11 billion TLS handshakes captured at a university network to identify fingerprinting patterns and flaws in tools like Signal, Lantern, and Snowflake. The library supports both mimicry and randomization strategies, and a public dataset is available [269], which helps developers to analyze authentic TLS fingerprints more effectively.

A recent method that provides very flexible censorship circumvention by using custom protocol specification files is Proteus [228]. Proteus addresses some of the shortcomings in Marionette [300], such as the fact that it may pose security risks because it runs user-specified plugins, making hosts vulnerable to malicious code executions. Additionally, it does not support multiple protocols or version upgrades, meaning both clients and proxies must update in sync whenever a change is needed [228]. Proteus comes with a full specification of the language itself, which can be used to write Proteus protocols. Its protocol language is clearly defined, even for developers without expertise. Proteus also supports multiple protocols and versions at the same time, so proxies can serve clients under different censorship conditions without requiring everyone to update simultaneously. Instead of expecting users to write protocol specifications themselves, the system relies on trusted groups (*e.g.,* the Tor Project) to create and share them.

7.6.2. GAN-based mimicry

In Section 7.2.2, we briefly introduced a Domain Fronting technique, named Meek [86], which obfuscates Tor traffic by encapsulating it into an HTTPS connection to a CDN. However, Meek is susceptible to side-channel attacks [311] that are capable of distinguishing Meek traffic from other HTTPS traffic [102]. To thwart such side-channel attacks, [102] trained a Generative Adversarial Network (GAN) on Meek and normal HTTPS traffic. The authors managed to increase the average False Positive Rate (FPR) from 0.183 to 0.834, and therefore, made it more difficult for censors to differentiate actual HTTPS traffic from GAN-based mimicry traffic.

Another GAN-based mimicry method, Voiceover, is introduced by [230]. It is a voice-based protocol tunnel designed to covertly transmit data by embedding it into audio that mimics the natural patterns of human conversation. It uses GANs to learn and reproduce the timing and cadence of real two-person voice conversations. Voiceover ensures that the voice stream includes realistic speech and silence intervals to make the traffic appear legitimate to network observers.

7.6.3. User behavior-based mimicry

Unlike programmable obfuscation techniques, which mimic protocol-level fingerprints and semantics, user behavior-based mimicry aims to mimic the more general browsing behavior of a user. OUStralopithecus [237] is one example of a tool that generates “replaceable overt content”. It creates many concurrent and realistic web browsing sessions to evade detection by traffic analysis systems and bot detectors. The sessions simulate user-like browsing actions, such as clicking links, switching tabs, navigating between sites, and idling. Compared to prior systems like Slitheen [279], Waterfall [272], and Slitheen++ [246], OUStralopithecus produces traffic that closely mimics human behavior, as confirmed by high bot-detection scores from Cloudflare’s Bot Management system [237].

7.7. Deep packet inspection evasion

The core purpose of Deep Packet Inspection (DPI) evasion is to disguise or modify traffic patterns, headers, or payloads in a way that prevents DPI systems from detecting or classifying the traffic as undesirable. Obfuscation, more broadly, refers to any technique used to hide or disguise communications. DPI evasion fits within this broader category as a specific set of methods aimed at bypassing DPI tools.

The following subsections present research on automated DPI evasion technique discovery and an example of HTTP Request Smuggling.

7.7.1. Evasion technique discovery

Recently, there have been several efforts to automate the discovery of Deep Packet Inspection (DPI) evasion strategies [44,45,161,195,207,263]. These approaches typically involve probing DPI systems by sending modified packets and determining which packets (and in what format) pass through the middlebox.

The authors of Geneva (Genetic Evasion) [45], presented an algorithm that automatically discovers new censorship evasion techniques. Geneva fuzz-tests DPI systems by dropping, tampering with, duplicating, or fragmenting packets, to observe the DPI system’s behavior. The authors conducted both in-lab and real-world measurements in China, India, and Kazakhstan and were able to identify several successful strategies to evade censorship. As already briefly mentioned in Section 6.2.4, several studies have enhanced Geneva with new features [44,161,195,196], and over time, it has been used to develop and test effective censorship evasion strategies against real-world censors. Since the evasion strategies are derived from measurements, this survey classifies Geneva primarily as a measurement method (see Table 2), although the resulting evasion techniques can be used to circumvent DPI-based censorship.

The authors of SymTCP [263] also created an automated way of discovering DPI evasion strategies by reverse-engineering the TCP state machines of middleboxes. SymTCP identified packets that are processed differently by the DPI and the end host with a white-box TCP implementation. They evaluated SymTCP against DPI systems like Snort and Zeek, but also against China’s GFW and were able to discover new evasion strategies.

Unlike [45,263] which predominantly focused on finding evasion techniques for the transport layer, [195] discovered new evasion strategies for the application-layer protocols HTTP and DNS. They extended Geneva to support application-layer fuzz-testing, which includes inserting new bytes, replacing bytes, or changing string cases in HTTP and DNS messages. Through real-world experiments targeting the DPI systems in China, India, and Kazakhstan, they discovered 86 new unique evasion strategies [195].

Similarly, [215] focused specifically on HTTP-based evasion and demonstrated that HTTP request smuggling can be used to circumvent censorship. Their approach creates HTTP requests containing an additional, smuggled request and conflicting length headers, which cause the censor and the destination server to parse the message differently. As a result, the censor overlooks the smuggled request, while the server processes it normally. They evaluated this technique against the censorship systems of China, Russia, and Iran, and found that it successfully evaded HTTP-based censorship in all three cases.

7.8. Covert tunneling

Covert tunneling is a technique used to transmit data secretly through a legitimate communication channel, which is often established using a cover protocol such as TLS. The proposed covert tunneling solutions range from more complex ones, such as game-based circumvention [199], to simpler ones, such as messenger-based circumvention [244] or email-based covert tunneling [287]. The following subsection reviews WebRTC-based, Wasm-based, video channel-based, video game-based tunneling and, lastly, two covert tunneling frameworks.

7.8.1. WebRTC-based tunneling

Protozoa [253] is a covert tunneling solution which uses point-to-point WebRTC tunnels to circumvent censorship. It modifies the WebRTC stack and replaces the video frame with an IP packet containing

the secret data. Like most other covert tunneling systems, Protozoa is application-agnostic, meaning that it can tunnel any type of IP-based traffic through WebRTC while ensuring that the statistical properties of the original WebRTC video stream remain intact. Protozoa has been evaluated in the wild and was successful in evading the censorship systems in China, Russia, and India.

A distributed WebRTC-based Censorship-Resistant Overlay Network (CRON) which relies upon Protozoa is presented in [252]. CRON addresses some of Protozoa's limitations, such as the limited resistance to Sybil attacks or traffic analysis. CRON uses small groups of trusted users, connected through chatrooms provided by WebRTC-enabled platforms. These chatrooms are used to establish covert circuits among multiple CRON nodes and users can join the network by inviting others via shareable links. The system enables proxy discovery and secure routing of information through CRON nodes.

A method that combines covert WebRTC-based tunneling and steganography is presented in [234]. The authors introduced Stegozoa, which, instead of tunneling traffic over WebRTC, steganographically embeds the information in the WebRTC stream. They showed that the system is resistant to traffic analysis and, therefore, not distinguishable from legitimate WebRTC traffic, while still achieving a reasonable throughput of 8.2–23.8 Kbps.

7.8.2. Wasm-based tunneling

The authors of [222] introduced a novel technique which leverages WebAssembly (Wasm) to create circumvention transports. The technique allows developers to write new circumvention methods once, package them as WebAssembly Transport Modules (WATMs), and distribute them without needing to update the entire app. The WATER library enables existing tools to load and use these WATMs, making it easier to deploy new anti-censorship methods across different software and devices. As part of the evaluation, WATER was able to achieve a throughput of 1830 Mbps, which is close to the baseline throughput of raw TCP traffic (2210 Mbps).

7.8.3. Video channel-based tunneling

One example of a video channel-based tunneling method is CovertCast [281], which tunnels information over live-streaming services, such as the one offered by YouTube. The method allows a user in an uncensored region to download data from websites, modulate the data into images and then live-stream the images to YouTube. A user in a censored region can then monitor the live-stream and demodulate the images (assuming that the live-streaming service is not blocked).

Similar to CovertCast, [274] introduced DeltaShaper, a censorship-resistant tool that enables covert communication over the Internet by tunneling TCP/IP traffic through videoconferencing streams, specifically Skype. Unlike CovertCast, which only supports limited data types (*e.g.,* video or web content), DeltaShaper allows users in censored regions to establish full TCP/IP connections to external proxy servers. It preserves the statistical properties of packets to avoid detection through traffic analysis and can tunnel protocols like FTP, SMTP, or HTTP.

7.8.4. Video game-based tunneling

Covertly tunneling solutions through video games were proposed by [199,200,209,290]. Rook [200] provides low latency communication over the servers of the online first-person shooter *Team Fortress 2*, by embedding the covert data into the network traffic. Similarly, Castle [290] implements a prototypical covert communication channel over three real-time strategy games: *0 A.D.*, *Aeons*, and *Conquerors*. The authors do not propose a solution to the bootstrapping problem as it relies on out-of-band channels to distribute game instance information. However, their relatively low maximum bandwidth of 320 bps suggests that Castle can be used for bootstrapping higher bandwidth channels. One of the most recent examples is Telepath [199], a *Minecraft*-based covert communication channel. It allows users to add censorship circumvention functionality to *Minecraft* in the form of a *Minecraft* modification ("mod"), which then embeds covert data into non-disruptive in-game messages. Telepath relies on a functional Minecraft client and server connection, although different types of traffic can be embedded into the covert channel. Table 3, therefore, classifies Telepath, as well as Rook and Castle, as not entirely application-agnostic.

Issues with video games might be the low pain threshold for censors to block a certain video game (*i.e.,* low incentives to keep the service available). Some censors are willing to accept high collateral damage to enforce restrictions. For instance, China has entirely blocked access to the majority of social media apps and CDNs, and has developed its own domestic alternatives. In this context, blocking access to a video game would likely not represent a particularly high barrier for censors. Despite the low threshold for blocking such circumvention channels, video games offer potential benefits in some contexts, such as relatively low operational expenses, especially when one can leverage existing game server infrastructure.

7.8.5. Tunneling frameworks

Turbo Tunnel [255] is a design pattern for censorship circumvention that adds dedicated session and reliability layer protocols, such as QUIC or SCTP, between the user's application and the obfuscation layer (*e.g.,* obfs4, meek, Snowflake). This inner layer offers sequence numbers, acknowledgments and retransmissions. Most other tunneling systems either depend on the transport's innate session protocol (*e.g.,* TCP) or build ad-hoc session management into each obfuscation protocol. Turbo Tunnel, in contrast, uses a transport-agnostic session layer.

The authors of [242] introduced Balboa, a link obfuscation framework that tunnels covert data by modifying network traffic from unmodified applications without altering their observable behavior. It identifies and replaces portions of known, pre-shared network traffic (*e.g.,* audio data) with compact pointers to that data. On the receiving side, it uses the shared model to reconstruct the original content.

8. The interdependency of internet censorship measurements and circumvention research

Internet censorship measurements and circumvention go hand in hand. Circumvention tools rely on censorship measurement data, and vice versa. Developers need to understand what is being blocked and how censorship works in different regions so that they can build better tools to circumvent censorship. Without up-to-date and accurate measurements, some circumvention tools may stop working.

At the same time, censorship measurements depend on some circumvention tools to obtain accurate data. For example, if researchers want to measure censorship in a country that enforces censorship, they often need to use circumvention tools to access blocked content and confirm whether it is being censored. Without circumvention tools, researchers may not be able to tell the difference between regular network problems and intentional censorship.

A number of studies in the last decade have conducted censorship measurements, which have led to the development of new circumvention strategies. One of the most promising is fuzzing-based measurements which have led to specific evasion strategies. In particular, the authors of Geneva [45] contributed significantly to the field of automated evasion technique discovery by developing a system that uses fuzzing to discover censorship evasion strategies. The system has since been extended and tested in real-world settings, including against filtering systems in China, India, and Kazakhstan [44,161,195,196]. More measurement studies that directly translate findings into new circumvention policies are presented in [26,29,42,46,169]. One notable example of a crowdsourcing-based censorship measurement system is C-Saw [169], which collects measurements from a large number of users who have given explicit consent. The system can dynamically select the fastest circumvention approach based on the type of observed filtering. As a result, it incentivizes censorship measurements through

circumvention by offering fast page load times of otherwise censored web pages.

The research resulting from GFWatch and GFW Report, including but not limited to [18,31,71,111,120,127,312], demonstrates the interdependency between Internet censorship measurement and circumvention research. GFWatch provides an empirical foundation for understanding how China's GFW operates in practice, what gets blocked, when, and under what technical conditions. For example, in [312] researchers uncovered a buffer over-read vulnerability in the GFW's DNS-injection subsystem, which caused certain middleboxes to leak up to 125 bytes of their own internal memory. Over a two-year measurement campaign, the authors reverse-engineered the injector's parsing logic and monitored how and when the GFW patched the vulnerability.

Similarly, in [111], the authors documented a new censorship mechanism deployed by the GFW as of April 2024. The GFW began decrypting the "Initial" QUIC packet at scale, inspecting it, and applying blocklists. They demonstrated that the GFW introduced new QUIC-filtering mechanisms alongside its existing censorship systems. They also found that the set of domains blocked over QUIC differs significantly from those censored via DNS or TLS, and that the QUIC-filtering process is imperfect, since even moderate traffic volume overwhelms it. They then proposed multiple circumvention strategies, such as adjusting the source port numbers of the QUIC connection, preceding QUIC handshake packets with arbitrary UDP datagrams, performing QUIC connection migrations, or fragmenting the QUIC Initial packet.

9. Remarks on reproducibility

Reproducibility in Internet censorship measurement research, or more generally in Internet measurement research, is a challenge. This can be true for several reasons: Firstly, censors may change their censorship policy at any time, making measurement results not always reproducible. Since censorship events often coincide with periods of political unrest or elections, as shown in [13–16], some measurements may only capture a temporary snapshot. Secondly, both censorship and circumvention methods are constantly evolving. For example, a researcher may discover a new way to evade censorship and publish their findings. Subsequently, the censors may update their systems to block the newly discovered circumvention technique, and further research that aims to reproduce the original findings may not be able to do so. Thirdly, inconsistent metrics (*e.g.,* the use of different success/failure thresholds, timeouts, or error classifications) can produce different results from the same experiment. Lastly, vantage points used in measurement studies are oftentimes not privately-owned endpoints, partly to avoid putting individuals at risk and partly due to performance and reliability considerations. However, measurements from enterprise-level infrastructure may not reflect the experiences of typical users on residential or mobile networks.

Therefore, longitudinal censorship studies, which capture the variability of censorship practices over prolonged periods of time [190], are particularly important for understanding how censorship evolves. Large-scale measurement platforms [37–39,125,126] contribute significantly by providing this type of insight.

Other recent research has focused on surveying Internet measurement studies, pointing out the limited reproducibility [313] and proposing solutions to improve the accuracy, generality, and reproducibility of web measurement studies by providing a reusable infrastructure that overcomes the limitations of current custom-built approaches [314]. One noteworthy program introduced in 2023 is the ACM Internet Measurement Conference's (IMC) replicability track [315], which supports submissions that aim to reproduce or replicate prior results from IMC or other conferences publishing Internet measurement research.

10. Discussion and conclusions

The reviewed studies on Internet censorship measurements reveal a strong focus on specific regions, particularly China, Russia, Iran, and India, due to their strict censorship regimes and geopolitical relevance. However, this focus has led to a potential imbalance, with some well-studied regions being repeatedly analyzed while others with censorship (*e.g.,* South America, Africa, or parts of Europe) receive little attention.

At the same time, it can be concluded that the focus of censorship measurements only partly aligns with actual censorship prevalence, as illustrated by countries like Cuba, where censorship levels are high [316] but research attention is almost nonexistent. There may also be a potential bias between actual censorship levels and research attention in African nations. For instance, although there is some research on Internet censorship and shutdowns in African countries, particularly on temporary restrictions around elections [17,317–319], there is very little quantitative research on Internet measurement. In response to this gap, there are initiatives to increase the coverage of Internet (censorship) measurements in Africa [36,320].

Moreover, ongoing conflicts and active war zones are very difficult to study. It is, for instance, challenging to differentiate between general Internet shutdowns and targeted censorship. For example, a longitudinal study [321] explicitly tracked Internet disruptions in Ukraine over a period of three years (2022–2025) and was able to remotely identify disruptions caused by the war; however, the underlying causes of the disruptions, such as power outages, destruction of infrastructure, or equipment seizure, could not always be unambiguously distinguished. The authors of [13] also conducted Internet measurements at the beginning of (and related to) the Ukraine–Russia war, but explicitly measured only geoblocking in Russia and not potential Internet disruptions in Ukraine.

Similarly, there are some studies that analyzed censorship around the Palestine–Israel conflict [322,323], but the number of studies is limited and data collection is challenging in regions where infrastructure boundaries are not clear. Performing actual Internet censorship measurements in such regions and obtaining unbiased results is very difficult. Our paper has analyzed many studies that still relied on volunteers and activists to deploy and maintain censorship measurement and circumvention infrastructure (*e.g.,* [3,37,169,170,259,292]), which is infeasible in active war zones.

The methodology of our paper may also cause under-representation of certain regions where strict Internet censorship is implemented, *i.e.,* regions where governments simply "switch off" the Internet instead of conducting targeted censorship. As a result, such regions provide a limited incentive for researchers to conduct measurements, as the primary observable outcome is a complete Internet disruption rather than targeted censorship.

In terms of methodology, Internet censorship measurement studies often focus on network protocols, such as HTTP, DNS, and TLS, because of their central role in web traffic. Longitudinal studies, especially those using data published by OONI, Censored Planet, and ICLab, are becoming increasingly relevant. These platforms provide open-access datasets and enable large-scale, longitudinal analysis of censorship worldwide.

Longitudinal studies are more common in censorship measurement research, whereas long-term testing of circumvention methods in the academic literature is less frequent. This is mainly due to practical challenges, such as maintaining infrastructure and adapting to new censorship tactics, but also because measuring the effectiveness of a circumvention method might put users at risk.

Although Internet censorship measurement studies commonly disclose the countries involved, Internet censorship circumvention studies, on the other hand, cannot disclose the real effectiveness of their circumvention tools on a per country basis, neither as part of short-term nor longitudinal studies. This would require attributing usage data of users experiencing censorship, which would expose users to risks.

The analysis of Internet censorship circumvention research has shown that there are two broad categories of techniques: routing-based and obfuscation-based. Routing-based methods can be further divided

into Proxying, Alibi Routing, and Connection Splitting. Obfuscation-based methods are Steganography, Protocol Mimicry, DPI Evasion, and Covert Tunneling. The analysis of protocols used shows a heavy reliance on TLS, while some more adaptable tools and frameworks enable dynamic protocol selection. The study also emphasizes the complexity of designing bootstrapping mechanisms and the challenges of achieving high throughput under censorship conditions. A taxonomy of tools reveals varying levels of application-agnosticism, with many being tailored to individual use-cases.

MTD systems, especially systems with ephemeral proxies like Snowflake [99], NetShuffle [192], and SpotProxy [220], offer strong resistance to proxy enumeration attacks by frequently changing IP addresses or proxy instances. Other proposals such as Alibi Routing [201] and DeTor [202] introduce verifiable geographic avoidance without requiring infrastructure changes. Connection Splitting techniques [226,241,250] leverage transport-layer features like QUIC connection migrations and TLS session resumptions to bypass traffic analysis. MTD approaches and frameworks for dynamic protocol selection (*e.g.,* [82,228,255]) indicate that adaptability is the key to fighting censors. These methods make enumeration and blocking much harder for adversaries.

Steganographic systems, such as MoneyMorph [249] and Discop [232], offer provably secure covert channels, while methods such as history covert channels [211,217] and SkypeLine [294] hide data in public content and voice traffic. Protocol mimicry tools like Marionette [300], Proteus [228], and uTLS [268] enable programmable obfuscation. We also reviewed covert tunneling techniques that hide data inside legitimate protocols. These include WebRTC-based systems like Protozoa [254], Wasm-based modules like WATER [222], and game-based tools like Telepath [199].

In circumvention research, there is a lot of focus on techniques that have only limited practicality and are not used on a wide scale. One such example is refraction networking, which has been the subject of at least 12 research papers in the past decade. Although there is a production deployment of refraction networking that has served thousands of users per month [261], building a truly robust and scalable system based on this approach would require the cooperation of multiple (ideally large) ISPs. In highly isolated regions where securing collaboration from ISPs may be very difficult, refraction networking is infeasible.

Another observable trend may be the overestimation of collateral damage to censors. In particular, niche video games used for censorship circumvention could be easily blocked without inflicting collateral, especially because some nation states are not hesitant to block even more widely-used services such as Gmail.

Despite the growing number of new circumvention techniques, their adoption remains low. Fundamental systems such as I2P and Tor, especially in combination with pluggable transports like Snowflake [98], obfs4 [90], and meek [100], continue to serve as essential tools for journalists and activists worldwide to circumvent state-level censorship.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests:

Thomas Grübl and Burkhard Stiller report that financial support was provided by (a) the University of Zürich UZH, Switzerland, (b) the Horizon Europe Framework Program's project Certify, Grant Agreement No. 101069471, funded by the Swiss State Secretariat for Education, Research, and Innovation SERI, under Contract No. 22.00165. Francisco Enguix reports that financial support was provided under the grant PID2021-123673OB-C31, PRE2022-101563, funded by MICIU/AEI/10.13039/501100011033 and by "ERDF A way of making Europe". If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work was partially supported by (a) the University of Zürich UZH, Switzerland, (b) the Horizon Europe Framework Program's project Certify, Grant Agreement No. 101069471, funded by the Swiss State Secretariat for Education, Research, and Innovation SERI, under Contract No. 22.00165, and (c) the grant PID2021-123673OB-C31, PRE2022-101563, funded by MICIU/AEI/10.13039/501100011033 and by "ERDF A way of making Europe".

We would like to thank Jürgen Bernard for his valuable feedback on the visualizations. We are also grateful to Philipp Winter for maintaining a curated bibliography of Internet censorship research—CensorBib, which is available at https://censorbib.nymity.ch/.

Special thanks to Gabriel Gegenhuber for sharing a template for the world map visualization, and to Eike Petersen for providing the template used for the boxplot.

We further acknowledge Roya Ensafi and her team for their significant contributions to the field, particularly through the development of Censored Planet and for their commitment to open-sourcing its data. Similarly, we want to acknowledge the Open Observatory of Network Interference (OONI) for their continued efforts to monitor and document Internet censorship worldwide, and for also making their data openly available to the research community. We are grateful to all researchers and organizations working to advance our understanding of Internet censorship and digital freedom.

Appendix A. Research areas related to internet censorship

See Table A.4.

Table A.4
Research area groups Ordered by the Overall number of Papers.

Category	Total	1995	1996	1998	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025
Internet Censorship	281	0	0	0	1	1	0	0	1	2	0	1	2	4	5	13	13	22	17	14	16	25	26	12	24	17	15	19	16	15
Internet	178	1	4	0	0	1	0	0	2	3	1	1	2	6	5	9	11	14	13	11	10	14	11	4	8	8	11	8	10	10
Media and Society	175	2	5	0	0	2	0	0	2	3	0	1	1	3	3	8	8	8	8	11	12	17	15	7	12	10	13	13	5	6
Country-Specific	120	2	4	0	0	1	0	0	1	3	0	0	1	2	4	1	3	11	10	8	7	10	8	5	9	4	5	8	8	5
Circumvention	95	0	0	0	0	0	0	0	0	1	0	0	0	1	1	5	4	6	6	6	8	9	7	5	7	5	4	6	7	7
Internet Protocols	93	1	1	1	0	1	0	0	0	2	0	0	0	3	2	6	7	5	7	3	8	9	3	5	6	6	3	3	5	6
Artificial Intelligence	82	0	1	0	0	0	0	0	0	0	0	1	0	1	1	9	3	5	2	8	4	5	10	4	8	7	3	2	3	5
Network Security	79	0	0	0	0	0	0	0	1	0	0	1	1	2	2	6	5	2	5	4	4	7	5	4	9	3	2	6	2	8
Internet Service Providers	76	0	0	0	0	1	0	0	1	1	1	1	0	0	1	6	3	9	6	4	4	4	4	3	8	4	3	5	3	4
Government	67	1	0	1	0	0	0	0	0	0	0	0	0	2	2	3	1	8	1	3	2	3	6	4	7	6	7	3	3	4
Routing	64	2	5	0	0	0	0	0	0	1	0	1	0	1	1	3	3	7	3	3	4	4	3	4	3	2	3	5	5	1
Data Privacy	64	0	0	0	0	0	0	0	1	0	0	0	1	1	1	1	3	6	4	5	7	8	0	5	3	6	1	7	1	3
Social Media	63	0	1	0	0	0	0	0	0	0	0	0	0	0	2	2	4	3	3	2	3	7	5	2	5	6	4	3	4	7
Search Engines	60	0	0	0	0	0	0	0	1	1	0	1	1	1	1	3	1	3	2	4	0	5	9	3	5	4	3	5	5	2
Traffic Analysis	56	0	0	0	0	0	0	0	0	0	0	1	0	0	3	2	3	5	3	5	1	6	5	2	5	3	4	4	3	1
Information Filtering	46	0	0	0	0	0	0	0	0	0	0	1	1	0	1	1	2	3	5	3	2	4	5	1	4	5	0	3	2	3
Measurements	45	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	0	3	2	4	4	5	2	4	2	4	7	2	3
Information Technology	45	0	0	0	0	1	0	0	0	1	0	0	0	2	1	1	3	4	2	1	3	4	1	0	8	4	3	2	4	0
Human Rights	37	1	4	0	0	0	0	0	1	0	1	0	0	2	2	2	3	3	3	2	1	2	1	1	2	3	1	2	0	0
Human Computer Interaction	35	0	2	0	0	0	0	0	1	0	0	0	0	1	1	1	0	3	2	2	3	7	3	0	3	2	1	2	0	1

Appendix B. Acronyms

Table B.5
List of acronyms.

Acronym	Definition	Acronym	Definition
AS	Autonomous System	LSA	Latent Semantic Analysis
ASCII	American Standard Code for Information Interchange	MF	MobilityFirst
AWS	Amazon Web Services	MIPv6	Mobile Internet Protocol version 6
BGP	Border Gateway Protocol	MTD	Moving Target Defense
CA	Certificate Authority	NDN	Named Data Networking
CDN	Content Delivery Network	OONI	Open Observatory of Network Interference
CoA	Care-of-Address	OpenWRT	Open Wireless Router
CRON	Censorship-Resistant Overlay Network	OSI	Open Systems Interconnection
CRS	Censorship Resistance Systems	P2P	Peer-to-Peer
CSV	Comma-Separated Values	PKI	Public Key Infrastructure
DNS	Domain Name System	PLD	Pay-Level Domain
DoT	DNS over TLS	PRNG	Pseudo Random Number Generator
DPI	Deep Packet Inspection	PTCL	Pakistan Telecommunication Company Ltd.
DSSS	Direct Sequence Spread Spectrum	RAD	Routing Around Decoys
ePDG	Evolved Packet Data Gateway	RSS	Really Simple Syndication
ESNI	Encrypted Server Name Indication	RTC	Real-Time Communication
EU	European Union	RTT	Round Trip Time
FIA	Future Internet Architecture	SCION	Scalability, Control, and Isolation On Next-Generation Networks
FQDN	Fully Qualified Domain Name	SCTP	Stream Control Transmission Protocol
FTP	File Transfer Protocol	SIM	Subscriber Identity Module
GAN	Generative Adversarial Network	SMTP	Simple Mail Transfer Protocol
GDPR	General Data Protection Regulation	SNI	Server Name Indication
GFW	Great Firewall (of China)	SQS	Simple Queue Service
GPS	Global Positioning System	SSH	Secure Shell
HTTP	Hypertext Transfer Protocol	TCP	Transmission Control Protocol
HTTPS	Hypertext Transfer Protocol Secure	TLD	Top-Level Domain
I2P	Invisible Internet Project	TLS	Transport Layer Security
ICMP	Internet Control Message Protocol	Tor	The Onion Routing
ICT	Information and Communication Technology	TPR	True Positive Rate
IDS	Intrusion Detection System	TTL	Time To Live
IKE	Internet Key Exchange	UDP	User Datagram Protocol
IODA	Internet Outage Detection and Analysis	URL	Uniform Resource Locator
IP	Internet Protocol	VoIP	Voice-over-IP
ISN	Initial Sequence Number	VPN	Virtual Private Network
ISP	Internet Service Provider	VPS	Virtual Private Server
IXP	Internet eXchange Point	Wasm	WebAssembly
KL	Kullback-Leibler	WATM	WebAssembly Transport Modules

Data availability

The data and supplementary materials for this research are available in a GitHub repository at https://github.com/thomasgruebl/internet-censorship-review

References

[1] R.S. Raman, A. Stoll, J. Dalek, Measuring the Deployment of Network Censorship Filters at Global Scale, in: Network and Distributed Systems Security (NDSS) Symposium 2020, 2020.
[2] R.S. Raman, P. Shenoy, K. Kohls, R. Ensafi, Censored planet: An Internet-wide, longitudinal censorship observatory, in: proceedings of the 2020 ACM SIGSAC conference on computer and communications security, 2020, pp. 49–66.
[3] N.P. Hoang, S. Doreen, M. Polychronakis, Measuring I2P censorship at a global scale, in: 9th USENIX Workshop on Free and Open Communications on the Internet (FOCI 19), 2019.
[4] The Citizen Lab, The citizen lab, 2025. https://citizenlab.ca.
[5] Amnesty International, Freedom of expression, 2025. https://www.amnesty.org/en/what-we-do/freedom-of-expression/.
[6] L. Taylor, What is data justice? The case for connecting digital rights and freedoms globally, Big Data Soc. 4 (2) (2017) 2053951717736335.
[7] R. Ramesh, A. Vyas, R. Ensafi, All of them claim to be the best: Multi-perspective study of VPN users and VPN providers, in: 32nd USENIX Security Symposium (USENIX Security 23), 2023, pp. 5773–5789.
[8] T. T. P. Inc, The Tor Project, 2025. https://www.torproject.org/.
[9] T. I. I. Project, The Invisible Internet Project, 2025. https://geti2p.net/en/.
[10] F. Project, Hyphanet, 2025. https://www.hyphanet.org/.
[11] F. House, Freedom in the World 2025: The Uphill Battle to Safeguard Rights, 2025. https://freedomhouse.org/sites/default/files/2025-03/FITW_World2025digitalN.pdf.
[12] M.S. Al-Zaman, M.M.S. Noman, Rise of Digital Authoritarianism? Exploring Global Motivations Behind Governmental Social Media Censorship, Journal of the European Institute for Communication and Culture 31 (4) (2024) 529–544.
[13] R. Ramesh, R.S. Raman, A. Virkud, A. Dirksen, A. Huremagic, D. Fifield, D. Rodenburg, R. Hynes, D. Madory, R. Ensafi, Network responses to Russia's invasion of Ukraine in 2022: a cautionary tale for Internet freedom, in: 32nd USENIX Security Symposium (USENIX Security 23), 2023, pp. 2581–2598.
[14] V. Ververis, T. Ermakova, M. Isaakidis, S. Basso, B. Fabian, S. Milan, Understanding Internet censorship in Europe: the case of Spain, in: Proceedings of the 13th ACM Web Science Conference 2021, 2021, pp. 319–328.
[15] R. Padmanabhan, A. Filastò, M. Xynou, R.S. Raman, K. Middleton, M. Zhang, D. Madory, M. Roberts, A. Dainotti, A multi-perspective view of Internet censorship in Myanmar, in: Proceedings of the ACM SIGCOMM 2021 Workshop on Free and Open Communications on the Internet, 2021, pp. 27–36.
[16] R.S. Raman, L. Evdokimov, E. Wurstrow, J.A. Halderman, R. Ensafi, Investigating large scale HTTPS interception in Kazakhstan, in: Proceedings of the ACM Internet Measurement Conference, 2020, pp. 125–132.
[17] T. Freyburg, L. Garbe, Blocking the bottleneck: Internet shutdowns and ownership at election times in sub-Saharan Africa, Int. J. Commun. 12 (2018) 3896–3916.
[18] M. Wu, J. Sippe, D. Sivakumar, J. Burg, P. Anderson, X. Wang, K. Bock, A. Houmansadr, D. Levin, E. Wustrow, How the Great Firewall of China detects and blocks fully encrypted traffic, in: 32nd USENIX Security Symposium (USENIX Security 23), 2023, pp. 2653–2670.
[19] K. Singh, G. Grover, V. Bansal, How India censors the web, in: Proceedings of the 12th ACM Conference on Web Science, 2020, pp. 21–28.
[20] A. McDonald, M. Bernhard, L. Valenta, B. VanderSloot, W. Scott, N. Sullivan, J.A. Halderman, R. Ensafi, 403 forbidden: a global view of CDN geoblocking, in: Proceedings of the 18th ACM Internet Measurement Conference, 2018, pp. 218–230.
[21] O. Farnan, A. Darer, J. Wright, Poisoning the well: Exploring the great firewall's poisoned DNS responses, in: Proceedings of the 2016 ACM on Workshop on Privacy in the Electronic Society, 2016, pp. 95–98.
[22] V. Ververis, G. Kargiotakis, A. Filastò, B. Fabian, A. Alexandros, Understanding Internet censorship policy: The case of Greece, in: 5th USENIX Workshop on Free and Open Communications on the Internet (FOCI 15), 2015.
[23] R.S. Raman, M. Wang, J. Dalek, J. Mayer, R. Ensafi, Network measurement methods for locating and examining censorship devices, in: Proceedings of the 18th International Conference on emerging Networking EXperiments and Technologies, 2022, pp. 18–34.
[24] D. Xue, R. Ramesh, L. Evdokimov, A. Viktorov, A. Jain, E. Wustrow, S. Basso, R. Ensafi, Throttling Twitter: an emerging censorship technique in Russia, in: Proceedings of the 21st ACM Internet Measurement Conference, 2021, pp. 435–443.
[25] T.K. Yadav, A. Sinha, D. Gosain, P.K. Sharma, S. Chakravarty, Where The Light Gets In: Analyzing Web Censorship Mechanisms in India, in: Proceedings of the Internet Measurement Conference 2018, 2018, pp. 252–264.
[26] F. Li, A. Razaghpanah, A.M. Kakhki, A.A. Niaki, D. Choffnes, P. Gill, A. Mislove, liberate, (n): a library for exposing (traffic-classification) rules and avoiding them efficiently, in: Proceedings of the 2017 Internet Measurement Conference, 2017, pp. 128–141.
[27] D. Katira, G. Grover, K. Singh, V. Bansal, CensorWatch: On the Implementation of Online Censorship in India, Free and Open Communications on the Internet (FOCI) (2023).

[28] A. Bhaskar, P. Pearce, Many Roads Lead To Rome: How Packet Headers Influence DNS Censorship Measurement, in: 31st USENIX Security Symposium (USENIX Security 22), 2022, pp. 449–464.
[29] J. Jermyn, N. Weaver, Autosonda: Discovering rules and triggers of censorship devices, in: 7th USENIX Workshop on Free and Open Communications on the Internet (FOCI 17), 2017.
[30] G. Aceto, A. Botta, A. Pescapé, M.F. Awan, T. Ahmad, S. Qaisar, Analyzing Internet censorship in Pakistan, in: 2016 IEEE 2nd International Forum on Research and Technologies for Society and Industry Leveraging a better tomorrow (RTSI), IEEE, 2016, pp. 1–6.
[31] N.P. Hoang, J. Dalek, M. Crete-Nishihata, N. Christin, V. Yegneswaran, M. Polychronakis, N. Feamster, GFWeb: Measuring the Great Firewall's Web Censorship at Scale, in: 33rd USENIX Security Symposium (USENIX Security 24), 2024.
[32] E. Tsai, R.S. Raman, A. Prakash, R. Ensafi, Modeling and Detecting Internet Censorship Events, in: Network and Distributed System Security (NDSS) Symposium, 2024.
[33] Z.S. Bischof, K. Pitcher, E. Carisimo, A. Meng, R.B. Nunes, R. Padmanabhan, M.E. Roberts, A.C. Snoeren, A. Dainotti, Destination Unreachable: Characterizing Internet Outages and Shutdowns, in: Proceedings of the ACM SIGCOMM 2023 Conference, 2023, pp. 608–621.
[34] P. Gill, M. Crete-Nishihata, J. Dalek, S. Goldberg, A. Senft, G. Wiseman, Characterizing web censorship worldwide: Another look at the opennet initiative data, ACM Transactions on the Web (TWEB) 9 (1) (2015) 1–29.
[35] A. Filasto, J. Appelbaum, OONI: Open Observatory of Network Interference, in: 2nd USENIX Workshop on Free and Open Communications on the Internet (FOCI 12), 2012.
[36] A.A. Niaki, S. Cho, Z. Weinberg, N.P. Hoang, A. Razaghpanah, N. Christin, P. Gill, Iclab: A global, longitudinal internet censorship measurement platform, in: 2020 IEEE Symposium on Security and Privacy (SP), IEEE, 2020, pp. 135–151.
[37] A. Filasto, J. Appelbaum, OONI: Open Observatory of Network Interference, 2012. https://ooni.org/.
[38] R.S. Raman, P. Shenoy, K. Kohls, R. Ensafi, Censored planet, 2020. https://censoredplanet.org/.
[39] A.A. Niaki, S. Cho, Z. Weinberg, N.P. Hoang, A. Razaghpanah, N. Christin, P. Gill, Iclab: A global, longitudinal internet censorship measurement platform, 2020. https://iclab.org/.
[40] C.L. Givens, Hidden forms of Censorship and their impact, Bookbird 47 (3) (2009) 22.
[41] P. Cook, C. Heilmann, Censorship and two types of self-censorship, 2010 6 (2).
[42] S. Nourin, V. Tran, X. Jiang, K. Bock, N. Feamster, N.P. Hoang, D. Levin, Measuring and evading Turkmenistan's Internet censorship: A case study in large-scale measurements of a low-penetration country, in: Proceedings of the ACM Web Conference 2023, 2023, pp. 1969–1979.
[43] A. Amich, B. Eshete, V. Yegneswaran, N.P. Hoang, DeResistor: Toward Detection-Resistant Probing for Evasion of Internet Censorship, in: 32nd USENIX Security Symposium (USENIX Security 23), 2023, pp. 2617–2633.
[44] K. Bock, G. Naval, K. Reese, D. Levin, Even censors have a backup: Examining China's double HTTPS censorship middleboxes, in: Proceedings of the ACM SIGCOMM 2021 Workshop on Free and Open Communications on the Internet, 2021, pp. 1–7.
[45] K. Bock, G. Hughey, X. Qiang, D. Levin, Geneva: Evolving censorship evasion strategies, in: Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security, 2019, pp. 2199–2214.
[46] Z. Wang, Y. Cao, Z. Qian, C. Song, S.V. Krishnamurthy, Your state is not mine: A closer look at evading stateful internet censorship, in: Proceedings of the 2017 Internet Measurement Conference, 2017, pp. 114–127.
[47] D. Xue, A. Ablove, R. Ramesh, G.K. Danciu, R. Ensafi, Bridging Barriers: A Survey of Challenges and Priorities in the Censorship Circumvention Landscape, in: 33rd USENIX Security Symposium (USENIX Security 24), 2024, pp. 2671–2688.
[48] M.C. Tschantz, S. Afroz, V. Paxson, et al., Sok: Towards grounding censorship circumvention in empiricism, in: 2016 IEEE Symposium on Security and Privacy (SP), IEEE, 2016, pp. 914–933.
[49] M. Wrana, D. Barradas, N. Asokan, The spectre of surveillance and censorship in future Internet architectures, in: Proceedings on Privacy Enhancing Technologies, 2025.
[50] A. Master, C. Garman, A Worldwide View of Nation-state Internet Censorship, Free and Open Communications on the Internet (FOCI) (2023).
[51] I. Makhdoom, M. Abolhasan, J. Lipman, A comprehensive survey of covert communication techniques, limitations and future challenges, Comput. Secur. 120 (2022) 102784.
[52] S. Wendzel, S. Volpert, S. Zillien, J. Lenz, P. Rünz, L. Caviglione, A Survey of Internet Censorship and its Measurement: Methodology, arXiv preprint arXiv:2502.14945, 2025.
[53] S. Khattak, M.T. Elahi, L. Simon, C.M. Swanson, S.J. Murdoch, I. Goldberg, SoK: Making Sense of Censorship Resistance Systems, in: Proceedings on Privacy Enhancing Technologies, 2016, pp. 37–61.
[54] G. Aceto, A. Pescapé, Internet censorship detection: A survey, Comput. Netw. 83 (2015) 381–421.
[55] F. Enguix, C. Carrascosa, J. Rincon, Exploring Federated Learning Tendencies Using a Semantic Keyword Clustering Approach, Information 15 (7) (2024) 379.
[56] H. Snyder, Literature review as a research methodology: An overview and guidelines, J. Bus. Res. 104 (2019) 333–339.
[57] B. Kitchenham, S. Charters, Guidelines for performing systematic literature reviews in software engineering, Tech. rep., School of Computer Science and Mathematics, Keele University, 2007.
[58] P. Winter, CensorBib: Selected Research Papers in Internet Censorship, 2024. https://censorbib.nymity.ch/.
[59] C. Wohlin, Guidelines for snowballing in systematic literature studies and a replication in software engineering, in: Proceedings of the 18th international conference on evaluation and assessment in software engineering, 2014, pp. 1–10.
[60] M. Antonakakis, T. April, M. Bailey, M. Bernhard, E. Bursztein, J. Cochran, Z. Durumeric, J.A. Halderman, L. Invernizzi, M. Kallitsis, et al., Understanding the mirai botnet, in: 26th USENIX security symposium (USENIX Security 17), 2017, pp. 1093–1110.
[61] Y. Mou, K. Wu, D. Atkin, Understanding the use of circumvention tools to bypass online censorship, New Media Soc. 18 (5) (2016) 837–856.
[62] J. Knockel, L. Ruan, M. Crete-Nishihata, An analysis of automatic image filtering on WeChat Moments, in: 8th USENIX Workshop on Free and Open Communications on the Internet (FOCI 18), 2018.
[63] Elsevier, Scopus, 2025. https://www.elsevier.com/products/scopus.
[64] all-mpnet-base-v2, 2025, https://huggingface.co/sentence-transformers/all-mpnet-base-v2.
[65] scikit-learn developers, AgglomerativeClustering, 2025. https://scikit-learn.org/stable/modules/generated/sklearn.cluster.AgglomerativeClustering.html.
[66] C. Dictionary, Censor, 2025. https://dictionary.cambridge.org/dictionary/english/censor.
[67] M.C. Tschantz, S. Afroz, S. Sajid, S.A. Qazi, M. Javed, V. Paxson, A bestiary of blocking: The motivations and modes behind website unavailability, in: 8th USENIX Workshop on Free and Open Communications on the Internet (FOCI 18), 2018.
[68] W. Rehan, S. Fischer, M. Rehan, Anatomizing the robustness of multichannel MAC protocols for WSNs: An evaluation under MAC oriented design issues impacting QoS, J. Netw. Comput. Appl. 121 (2018) 89–118, https://doi.org/10.1016/j.jnca.2018.06.013, https://www.sciencedirect.com/science/article/pii/S1084804518302212.
[69] Z. Weinberg, D. Barradas, N. Christin, Chinese wall or Swiss cheese? Keyword filtering in the great firewall of China, in: Proceedings of the Web Conference 2021, 2021, pp. 472–483.
[70] R. Ramesh, R.S. Raman, M. Bernhard, V. Ongkowijaya, L. Evdokimov, A. Edmundson, S. Sprecher, M. Ikram, R. Ensafi, Decentralized control: A case study of Russia, in: Network and Distributed Systems Security (NDSS) Symposium 2020, 2020.
[71] B. Alice, J.B. Carol, A. Houmansadr, How China detects and blocks shadowsocks, in: Proceedings of the ACM Internet Measurement Conference, 2020, pp. 111–124.
[72] C. Dictionary, Collateral Damage, 2025. https://dictionary.cambridge.org/dictionary/english/collateral-damage.
[73] O. Initiative, OpenNet, 2002. https://opennet.net/.
[74] X. Xu, Z.M. Mao, J.A. Halderman, Internet censorship in China: Where does the filtering occur? in: Passive and Active Measurement: 12th International Conference, PAM 2011, (20 March 2011). Proceedings 12, Springer, Atlanta, GA, USA, 2011, pp. 133–142.
[75] N. Aase, J.R. Crandall, Á. Díaz, J. Knockel, J.O. Molinero, J. Saia, D.S. Wallach, T. Zhu, Whiskey, Weed, and Wukan on the World Wide Web: On Measuring Censors' Resources and Motivations, FOCI, 2012.
[76] R. Clayton, S.J. Murdoch, R.N. Watson, Ignoring the great firewall of China, in: International workshop on privacy enhancing technologies, Springer, 2006, pp. 20–35.
[77] N.P. Hoang, A.A. Niaki, J. Dalek, J. Knockel, P. Lin, B. Marczak, M. Crete-Nishihata, P. Gill, M. Polychronakis, How Great is the Great Firewall? Measuring China's DNS Censorship, in: 30th USENIX Security Symposium (USENIX Security 21), 2021, pp. 3381–3398.
[78] J.R. Crandall, D. Zinn, M. Byrd, E.T. Barr, R. East, ConceptDoppler: a weather tracker for internet censorship, CCS 7 (2007) 352–365.
[79] R. Deibert, J. Palfrey, R. Rohozinski, J. Zittrain, Access controlled: The shaping of power, rights, and rule in cyberspace, the, MIT Press, 2010.
[80] I. Clarke, et al., A distributed decentralised information storage and retrieval system, Division of Informatics, University of Edinburgh, 1999, pp. 1–45.
[81] N. Feamster, M. Balazinska, G. Harfst, H. Balakrishnan, D. Karger, Infranet: Circumventing web censorship and surveillance, in: 11th USENIX Security Symposium (USENIX Security 02), 2002.
[82] P. Vines, S. McKay, J. Jenter, S. Krishnaswamy, Communication Breakdown: Modularizing Application Tunneling for Signaling Around Censorship, in: Proceedings on Privacy Enhancing Technologies, 2024, pp. 465–477.
[83] A. Houmansadr, G.T. Nguyen, M. Caesar, N. Borisov, Cirripede: Circumvention infrastructure using router redirection with plausible deniability, in: Proceedings of the 18th ACM conference on Computer and communications security, 2011, pp. 187–200.
[84] D. Fifield, G. Nakibly, D. Boneh, Oss: Using online scanning services for censorship circumvention, in: International Symposium on Privacy Enhancing Technologies Symposium, Springer, 2013, pp. 185–204.
[85] M. Nasr, S. Farhang, A. Houmansadr, J. Grossklags, Enemy At the Gateways: Censorship-Resilient Proxy Distribution Using Game Theory, in: Network and Distributed Systems Security (NDSS) Symposium 2019, 2019.
[86] D. Fifield, C. Lan, R. Hynes, P. Wegmann, V. Paxson, Blocking-resistant communication through domain fronting, in: Proceedings on Privacy Enhancing Technologies, 2015.
[87] Microsoft, Securing our approach to domain fronting within Azure, 2021. https://www.microsoft.com/en-us/security/blog/2021/03/26/securing-our-approach-to-domain-fronting-within-azure/.
[88] S. Gallagher, Google disables "domain fronting" capability used to evade censors, 2018. https://arstechnica.com/information-technology/2018/04/google-disables-domain-fronting-capability-used-to-evade-censors/.

[89] P. Lincoln, I. Mason, P.A. Porras, V. Yegneswaran, Z. Weinberg, J. Massar, W.A. Simpson, P. Vixie, D. Boneh, Bootstrapping Communications into an Anti-Censorship System, Free and Open Communications on the Internet (FOCI) (2012).
[90] Y. Angel, obfs4, 2014. https://github.com/Yawning/obfs4/tree/master.
[91] D.M. Goldschlag, M.G. Reed, P.F. Syverson, Hiding routing information, in: International workshop on information hiding, Springer, 1996, pp. 137–150.
[92] C. Adams, Introduction to privacy enhancing technologies: a classification-based approach to understanding PETs, Springer Nature, 2021.
[93] M.G. Reed, P.F. Syverson, D.M. Goldschlag, Anonymous connections and onion routing, IEEE J. Sel. Areas Commun. 16 (4) (1998) 482–494.
[94] M.G. Reed, P.F. Syverson, D.M. Goldschlag, Proxies for anonymous routing, in: Proceedings 12th Annual Computer Security Applications Conference, IEEE, 1996, pp. 95–104.
[95] M.G. Reed, P.F. Syverson, D.M. Goldschlag, Protocols using anonymous connections: Mobile applications, in: Security Protocols: 5th International Workshop, (7 April 1997) Proceedings 5, Springer, Paris, France, 1998, pp. 13–23.
[96] P.F. Syverson, M.G. Reed, D.M. Goldschlag, Onion Routing access configurations, in: Proceedings DARPA Information Survivability Conference and Exposition. DISCEX'00, vol. 1, IEEE, 2000, pp. 34–40.
[97] R. Dingledine, N. Mathewson, P.F. Syverson, et al., Tor: The second-generation onion router, in: USENIX security symposium, vol. 4, 2004, pp. 303–320.
[98] C. Bocovich, A. Breault, D.F. Serene, X. Wang, Snowflake, 2024. https://snowflake.torproject.org/.
[99] C. Bocovich, A. Breault, D. Fifield, X.W. Serene, Snowflake, a censorship circumvention system using temporary WebRTC proxies, in: 33rd USENIX Security Symposium (USENIX Security 24), 2024, pp. 2635–2652.
[100] D. Fifield, C. Lan, R. Hynes, P. Wegmann, V. Paxson, meek, 2015. https://gitlab.torproject.org/legacy/trac/-/wikis/doc/meek.
[101] S. Sheffey, F. Aderholdt, Improving meek with adversarial techniques, in: 9th USENIX Workshop on Free and Open Communications on the Internet (FOCI 19), 2019.
[102] S. Sheffey, F. Aderholdt, Improving meek with adversarial techniques, 2019. https://github.com/starfys/packet_captor_sakura.
[103] G. shelikhoo, Hiding in plain sight: Introducing WebTunnel, 2024. https://blog.torproject.org/introducing-webtunnel-evading-censorship-by-hiding-in-plain-sight.
[104] I.I. Project, Garlic Routing and "Garlic", Terminology (2025) https://geti2p.net/en/docs/how/garlic-routing.
[105] J.P. Timpanaro, C. Isabelle, F. Olivier, Monitoring the I2P network (Ph.D. thesis), Centre de Recherche INRIA Nancy, 2011.
[106] S. Nourin, E. Rye, K. Bock, N.P. Hoang, D. Levin, Is Nobody There? Good! Globally Measuring Connection Tampering without Responsive Endhosts, in: 2025 IEEE Symposium on Security and Privacy (SP), IEEE, 2025, pp. 1400–1418.
[107] K. Bock, G. Hughey, X. Qiang, D. Levin, Geneva: Evolving Censorship Evasion, 2019. https://geneva.cs.umd.edu/.
[108] S. Nourin, E. Rye, K. Bock, N.P. Hoang, D. Levin, Exposing and Circumventing SNI-based QUIC Censorship of the Great Firewall of China, 2025. https://censorship.ai.
[109] N. Niere, F. Lange, N. Heitmann, J. Somorovsky, Encrypted Client Hello (ECH) in censorship circumvention, in: Free and Open Communications on the Internet, 2025, pp. 64–73.
[110] F. Lange, N. Niere, J. von Niessen, D. Suermann, N. Heitmann, J. Somorovsky, I(ra)nconsistencies: novel insights into Iran's censorship, in: Free and Open Communications on the Internet, 2025, pp. 7–12.
[111] A. Zohaib, Q. Zao, J. Sippe, A. Alaraj, A. Houmansadr, Z. Durumeric, E. Wustrow, Exposing and Circumventing SNI-based QUIC Censorship of the Great Firewall of China, in: 34th USENIX Security Symposium (USENIX Security 25), 2025, pp. 783–802.
[112] A. Zohaib, Q. Zao, J. Sippe, A. Alaraj, A. Houmansadr, Z. Durumeric, E. Wustrow, Exposing and Circumventing SNI-based QUIC Censorship of the Great Firewall of China, 2025. https://github.com/gfw-report/usenixsecurity25-quic-sni.
[113] D. Xue, A. Huremagic, W. Wang, R.S. Raman, R. Ensafi, Fingerprinting Deep Packet Inspection Devices by Their Ambiguities, in: Proceedings of the 2025 ACM SIGSAC Conference on Computer and Communications Security, 2025, pp. 3945–3959.
[114] D. Xue, A. Huremagic, W. Wang, R.S. Raman, R. Ensafi, dMAP, 2025. https://github.com/censoredplanet/CenDPI.
[115] D. Saha, A. Anwar, A. Mueen, Internet Censorship through the Lens of Time Series Analysis, in: Proceedings of the 17th ACM Web Science Conference 2025, 2025, pp. 534–539.
[116] J. Tai, K.N. Sengottuvelavan, P. Whiting, N.P. Hoang, IRBlock: A Large-Scale Measurement Study of the Great Firewall of Iran, in: 34th USENIX Security Symposium (USENIX Security 25), 2025, pp. 705–722.
[117] J. Tai, K.N. Sengottuvelavan, P. Whiting, N.P. Hoang, IRBlock: A Large-Scale Measurement Study of the Great Firewall of Iran (USENIX Security '25) - Artifacts, 2025. https://zenodo.org/records/15572895.
[118] F. Lipphardt, A. Feldmann, D. Gosain, Can You Hear Me? A First Study of VoIP Censorship Techniques in Saudi Arabia and the UAE, in: 2025 IEEE 10th European Symposium on Security and Privacy (EuroS&P), IEEE, 2025, pp. 720–736.
[119] F. Lipphardt, A. Feldmann, D. Gosain, Can You Hear Me? A First Study of VoIP Censorship Techniques in Saudi Arabia and the UAE, 2025. https://github.com/Fredddi43/eurosp_voip_censorship_artifacts.
[120] M. Wu, A. Zohaib, Z. Durumeric, A. Houmansadr, E. Wustrow, A Wall Behind A Wall: Emerging Regional Censorship in China, in: 2025 IEEE Symposium on Security and Privacy (SP), IEEE, 2025, pp. 1363–1380.
[121] M. Wu, A. Zohaib, Z. Durumeric, A. Houmansadr, E. Wustrow, A Wall Behind A Wall: Emerging Regional Censorship in China, 2025. https://gfw.report/publications/sp25/en/.
[122] A. Alaraj, E. Wustrow, Proxies as Sensors: Measuring Censorship of Refraction Networking in Iran, in: Proceedings of the 20th ACM Asia Conference on Computer and Communications Security, 2025, pp. 759–772.
[123] P. Inc, Psiphon, 2024. https://psiphon.ca/.
[124] A. Bhaskar, P. Pearce, Understanding Routing-Induced Censorship Changes Globally, in: Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security, 2024, pp. 437–451.
[125] N.P. Hoang, A.A. Niaki, J. Dalek, J. Knockel, P. Lin, B. Marczak, M. Crete-Nishihata, P. Gill, M. Polychronakis, GFWatch, 2021. https://gfwatch.org/.
[126] N.P. Hoang, J. Dalek, M. Crete-Nishihata, N. Christin, V. Yegneswaran, M. Polychronakis, N. Feamster, GFWeb, 2024. https://gfweb.ca/.
[127] E. Wedwards, et al., Bleeding Wall: A Hematologic Examination on the Great Firewall, Free and Open Communications on the Internet (FOCI) (2024).
[128] D. Gosain, K. Singh, R. Sharma, J.S. Babu, S. Chakravaty, Out in the Open: On the Implementation of Mobile App Filtering in India, in: International Conference on Passive and Active Network Measurement, Springer, 2024, pp. 19–36.
[129] G.K. Gegenhuber, P.É. Frenzel, E. Weippl, Why ET Can't Phone Home: A Global View on IP-based Geoblocking at VoWiFi, in: Proceedings of the 22nd Annual International Conference on Mobile Systems, Applications and Services, 2024, pp. 183–195.
[130] G.K. Gegenhuber, P.É. Frenzel, E. Weippl, scanywhere, 2024. https://github.com/sbaresearch/scanywhere.
[131] A. Amich, B. Eshete, V. Yegneswaran, N.P. Hoang, DeResistor, 2023. https://github.com/um-dsp/DeResistor.
[132] S.P. Duncan, H. Chen, Detecting network-based Internet censorship via latent feature representation learning, Comput. Secur. 128 (2023) 103138.
[133] S.P. Duncan, H. Chen, Censored Planet Processes and Models for Machine Learning, 2023. https://github.com/fathershawn/cp_learning.
[134] M. Wu, J. Sippe, D. Sivakumar, J. Burg, P. Anderson, X. Wang, K. Bock, A. Houmansadr, D. Levin, E. Wustrow, How the great firewall of china detects and blocks fully encrypted traffic, 2023. https://github.com/gfw-report/usenixsecurity23-artifact.
[135] R. Ramesh, R.S. Raman, A. Virkud, A. Dirksen, A. Huremagic, D. Fifield, D. Rodenburg, R. Hynes, D. Madory, R. Ensafi, GeoInspector, 2023. https://github.com/censoredplanet/geoinspector.
[136] D. Katira, G. Grover, K. Singh, V. Bansal, CensorWatch, 2023. https://cis-india.github.io/censorwatch/data.html.
[137] V. Ververis, L. Lasota, T. Ermakova, B. Fabian, Website blocking in the European Union: Network interference from the perspective of Open Internet, Policy Internet 16 (1) (2023) 121–148.
[138] CAIDA, Georgia Institute of Technology, IODA: Internet Outage Detection and Analysis, 2024. https://ioda.caida.org/ioda/dashboard.
[139] AccessNow, #KeepItOn: fighting internet shutdowns around the world, 2024. https://www.accessnow.org/campaign/keepiton/.
[140] Z.S. Bischof, K. Pitcher, E. Carisimo, A. Meng, R.B. Nunes, R. Padmanabhan, M.E. Roberts, A.C. Snoeren, A. Dainotti, Destination Unreachable, 2023. https://github.com/InetIntel/internet_outages.
[141] J. Brown, X. Jiang, V. Tran, A.N. Bhagoji, N.P. Hoang, N. Feamster, P. Mittal, V. Yegneswaran, Augmenting rule-based dns censorship detection at scale with machine learning, in: Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, 2023, pp. 3750–3761.
[142] J. Brown, X. Jiang, V. Tran, A.N. Bhagoji, N.P. Hoang, N. Feamster, P. Mittal, V. Yegneswaran, Automated DNS Censorship using Machine Learning, 2023. https://github.com/noise-lab/automated-dns-censorship.
[143] A. Ortwein, K. Bock, D. Levin, Towards a Comprehensive Understanding of Russian Transit Censorship, in: Workshop on Free and Open Communication on the Internet (FOCI), 2023.
[144] R.S. Raman, L.-H. Merino, K. Bock, M. Fayed, D. Levin, N. Sullivan, L. Valenta, Global, passive detection of connection tampering, in: Proceedings of the ACM SIGCOMM 2023 Conference, 2023, pp. 622–636.
[145] S. Nourin, V. Tran, X. Jiang, K. Bock, N. Feamster, N.P. Hoang, D. Levin, Turkmenistan Censorship, 2023. https://zenodo.org/records/7631411.
[146] R. Kumar, A. Virkud, R.S. Raman, A. Prakash, R. Ensafi, A large-scale investigation into geodifferences in mobile apps, in: 31st USENIX Security Symposium (USENIX Security 22), 2022, pp. 1203–1220.
[147] R. Kumar, A. Virkud, R.S. Raman, A. Prakash, R. Ensafi, geodiff-app, 2022. https://github.com/censoredplanet/geodiff-app.
[148] W. Scott, T. Anderson, T. Kohno, A. Krishnamurthy, Satellite: Mapping the internet's stars, 2016. https://github.com/UWNetworksLab/satellite.
[149] P. Pearce, B. Jones, F. Li, R. Ensafi, N. Feamster, N. Weaver, V. Paxson, Satellite/iris: Global detection of dns-layer disruption, 2017. https://censoredplanet.org/projects/satellite.
[150] D. Xue, B. Mixon-Baca, A.A. ValdikSS, B. Kujath, J.R. Crandall, R. Ensafi, TSPU: Russia's decentralized censorship system, in: Proceedings of the 22nd ACM Internet Measurement Conference, 2022, pp. 179–194.
[151] N.P. Hoang, M. Polychronakis, P. Gill, Measuring the accessibility of domain name encryption and its impact on Internet filtering, in: International Conference on Passive and Active Network Measurement, Springer, 2022, pp. 518–536.
[152] R.S. Raman, M. Wang, Centrace, 2022. https://github.com/censoredplanet/centrace.
[153] R.S. Raman, M. Wang, Cenfuzz, 2022. https://github.com/censoredplanet/CenFuzz.

[154] K. Elmenhorst, B. Schütz, N. Aschenbruck, S. Basso, Web censorship measurements of HTTP/3 over QUIC, in: Proceedings of the 21st ACM Internet Measurement Conference, 2021, pp. 276–282.
[155] S. Basso, Measuring DoT/DoH blocking using OONI probe: a preliminary study, in: NDSS DNS Privacy Workshop, 2021.
[156] B. Alice, J.B. Carol, A. Houmansadr, How China detects and blocks shadowsocks, 2020. https://gfw.report/publications/imc20/en/.
[157] R.S. Raman, A. Stoll, J. Dalek, FilterMap, 2020. https://censoredplanet.org/filtermap.
[158] A.A. Niaki, N.P. Hoang, P. Gill, A. Houmansadr, et al., Triplet Censors: Demystifying Great Firewall's DNS Censorship Behavior, in: 10th USENIX Workshop on Free and Open Communications on the Internet (FOCI 20), 2020.
[159] A.A. Niaki, N.P. Hoang, P. Gill, A. Houmansadr, et al., Triplet Censors: Demystifying Great Firewall's DNS Censorship Behavior, 2020. https://gfw.report/publications/foci20_dns/en/.
[160] F. Alharbi, M. Faloutsos, N. Abu-Ghazaleh, Opening digital borders cautiously yet decisively: Digital filtering in Saudi Arabia, in: 10th USENIX Workshop on Free and Open Communications on the Internet (FOCI 20), 2020.
[161] K. Bock, Y. Fax, K. Reese, J. Singh, D. Levin, 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), 2020.
[162] Z. Chai, A. Ghafari, A. Houmansadr, On the Importance of Encrypted-SNI (ESNI) to Censorship Circumvention, in: 9th USENIX Workshop on Free and Open Communications on the Internet (FOCI 19), 2019.
[163] R. Deibert, J. Oliver, A. Senft, Censors get smart: Evidence from Psiphon in Iran, Review of Policy Research 36 (3) (2019) 341–356.
[164] A. Dunna, C. O'Brien, P. Gill, Analyzing China's blocking of unpublished tor bridges, in: 8th USENIX Workshop on Free and Open Communications on the Internet (FOCI 18), 2018.
[165] P. Winter, S. Lindskog, How the great firewall of China is blocking Tor, 2012. https://www.usenix.org/sites/default/files/conference/protected-files/winter_foci12_slides.pdf.
[166] A. Dunna, C. O'Brien, P. Gill, Analyzing China's Blocking of Unpublished Tor Bridges, 2018. https://calipr.cs.umass.edu/projects/china-tor-arun.html.
[167] B. VanderSloot, A. McDonald, W. Scott, J.A. Halderman, R. Ensafi, Quack: Scalable Remote Measurement of Application-Layer Censorship, in: 27th USENIX Security Symposium (USENIX Security 18), 2018, pp. 187–202.
[168] B. VanderSloot, A. McDonald, W. Scott, J.A. Halderman, R. Ensafi, Quack, 2018. https://censoredplanet.org/projects/hyperquack.
[169] A. Nisar, A. Kashaf, I.A. Qazi, Z.A. Uzmi, Incentivizing censorship measurements via circumvention, in: Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication, 2018, pp. 533–546.
[170] V. Ververis, M. Isaakidis, C. Loizidou, B. Fabian, Internet censorship capabilities in Cyprus: An investigation of online gambling blocklisting, in: E-Democracy–Privacy-Preserving, Secure, Intelligent E-Government Services: 7th International Conference, Springer, Athens, Greece, 2017, pp. 136–149.
[171] V. Ververis, M. Isaakidis, C. Loizidou, B. Fabian, bet2512, 2017. https://github.com/hack66/bet2512.
[172] F. Li, A. Razaghpanah, A.M. Kakhki, A.A. Niaki, D. Choffnes, P. Gill, A. Mislove, liberate, (n): a library for exposing (traffic-classification) rules and avoiding them efficiently, 2017. https://dd.meddle.mobi/liberate.html.
[173] A. Darer, O. Farnan, J. Wright, FilteredWeb: A framework for the automated search-based discovery of blocked URLs, in: 2017 Network Traffic Measurement and Analysis Conference (TMA), IEEE, 2017, pp. 1–9.
[174] P. Pearce, R. Ensafi, F. Li, N. Feamster, V. Paxson, Augur: Internet-wide detection of connectivity disruptions, in: 2017 IEEE Symposium on Security and Privacy (SP), IEEE, 2017, pp. 427–443.
[175] P. Pearce, B. Jones, F. Li, R. Ensafi, N. Feamster, N. Weaver, V. Paxson, Global measurement of dns manipulation, in: 26th USENIX Security Symposium (USENIX Security 17), 2017, pp. 307–323.
[176] S. Cho, R. Nithyanand, A. Razaghpanah, P. Gill, A churn for the better: Localizing censorship using network-level path churn and network tomography, in: Proceedings of the 13th International Conference on emerging networking experiments and technologies, 2017, pp. 81–87.
[177] D. Gosain, A. Agarwal, S. Shekhawat, H.B. Acharya, S. Chakravarty, Mending wall: On the implementation of censorship in India, in: International Conference on Security and Privacy in Communication Systems, Springer, 2017, pp. 418–437.
[178] Z. Wang, Y. Cao, Z. Qian, C. Song, S.V. Krishnamurthy, INTANG, 2017. https://github.com/seclab-ucr/INTANG.
[179] W. Scott, T. Anderson, T. Kohno, A. Krishnamurthy, Satellite: Joint analysis of cdns and network-level interference, in: 2016 USENIX Annual Technical Conference (USENIX ATC, vol. 16, 2016, pp. 195–208.
[180] S. Burnett, N. Feamster, Encore: Lightweight measurement of web censorship with cross-origin requests, in: Proceedings of the 2015 ACM conference on special interest group on data communication, 2015, pp. 653–667.
[181] R. Ensafi, P. Winter, A. Mueen, J.R. Crandall, Analyzing the Great Firewall of China over space and time, in: Proceedings on Privacy Enhancing Technologies, 2015, pp. 61–76.
[182] G. Aceto, A. Botta, A. Pescapè, N. Feamster, M.F. Awan, T. Ahmad, S. Qaisar, Monitoring Internet censorship with UBICA, in: Traffic Monitoring and Analysis: 7th International Workshop, Springer, 2015, pp. 143–157.
[183] B. Jones, R. Ensafi, N. Feamster, V. Paxson, N. Weaver, Ethical concerns for censorship measurement, in: Proceedings of the 2015 ACM SIGCOMM Workshop on Ethics in Networked Systems Research, 2015, pp. 17–19.
[184] B. Jones, N. Feamster, Can Censorship Measurements Be Safe (r)? in: Proceedings of the 14th ACM Workshop on Hot Topics in Networks, 2015, pp. 1–7.
[185] A. Narayanan, B. Zevenbergen, No encore for encore? Ethical questions for web-based censorship measurement, in: Ethical Questions for Web-Based Censorship Measurement, 2015, pp. 1–21.
[186] D. Dittrich, E. Kenneally, The Menlo Report: Ethical Principles Guiding Information and Communication Technology Research, Tech. rep., U.S. Department of Homeland Security, 2012.
[187] U.S, The Belmont Report: Ethical Principles and Guidelines for the Protection of Human Subjects of Research, Tech. rep., U.S. Department of Health, Education, and Welfare, 1979.
[188] F. House, Internet Freedom Scores, 2024. https://freedomhouse.org/countries/freedom-net/scores.
[189] C. Angelopoulos, Filtering the Internet for copyrighted content in Europe, IRIS plus (Supplement to IRIS-Legal Observations of the European Audiovisual Observatory) 4 (2009) 04–2012.
[190] E.J. Caruana, M. Roman, J. Hernández-Sánchez, P. Solli, Longitudinal studies, J. Thorac. Dis. 7 (11) (2015) 537–540.
[191] P.K. Sharma, D. Xue, A. Ortwein, C. Bocovich, R. Ensafi, et al., CenPush: Blocking-Resistant Control Channel Using Push Notifications, in: Proceedings on Privacy Enhancing Technologies, 2025.
[192] P.T.J. Kon, A. Gattani, D. Saharia, T. Cao, D. Barradas, A. Chen, M. Sherr, B.E. Ujcich, NetShuffle: Circumventing Censorship with Shuffle Proxies at the Edge, in: IEEE Symposium on Security and Privacy (SP), 2024, pp. 36.
[193] M. Pu, A. Wang, A. Chang, K. Quan, Y.W. Zhou, Exploring Amazon Simple Queue Service (SQS) for censorship circumvention, in: Free and Open Communications on the Internet, 2024, pp. 22–26.
[194] D. Xue, R. Ensafi, The use of push notification in censorship circumvention, Free and Open Communications on the Internet (1) (2023) 22–32.
[195] M. Harrity, K. Bock, F. Sell, D. Levin, GET/out: Automated discovery of Application-Layer censorship evasion strategies, in: 31st USENIX Security Symposium (USENIX Security 22), 2022, pp. 465–483.
[196] K. Bock, G. Hughey, L.-H. Merino, T. Arya, D. Liscinsky, R. Pogosian, D. 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, 2020, pp. 586–598.
[197] R. Networks, RiseupVPN, 2025. https://riseup.net/vpn.
[198] A. Vilalonga, J.S. Resende, H. Domingos, Looking at the clouds: leveraging pub/sub cloud services for censorship-resistant rendezvous channels, in: Free and Open Communications on the Internet, 2024, pp. 27–33.
[199] Z. Sun, V. Shmatikov, Telepath: A minecraft-based covert communication system, in: 2023 IEEE Symposium on Security and Privacy (SP), IEEE, 2023, pp. 2223–2237.
[200] P. Vines, T. Kohno, Rook: Using video games as a low-bandwidth censorship resistant communication platform, in: Proceedings of the 14th ACM Workshop on Privacy in the Electronic Society, 2015, pp. 75–84.
[201] D. Levin, Y. Lee, L. Valenta, Z. Li, V. Lai, C. Lumezanu, N. Spring, B. Bhattacharjee, Alibi routing, in: ACM Conference on Special Interest Group on Data Communication (SIGCOMM), 2015, pp. 611–624.
[202] Z. Li, S. Herwig, D. Levin, DeTor: Provably Avoiding Geographic Regions in Tor, in: 26th USENIX Security Symposium, 2017, pp. 343–359.
[203] J. Ding, K. Chen, Y. Wang, N. Zhao, W. Zhang, N. Yu, Discop: Provably secure steganography in practice based on "distribution copies", in: 2023 IEEE Symposium on Security and Privacy (SP), IEEE, 2023, pp. 2238–2255.
[204] P.K. Sharma, D. Xue, A. Ortwein, C. Bocovich, R. Ensafi, et al., CenPush: Blocking-Resistant Control Channel Using Push Notifications, 2025. https://people.torproject.org/cohosh/push-notifications.html.
[205] R. Wails, R. Jansen, A. Johnson, M. Sherr, Censorship evasion with unidentified protocol generation, in: 34th USENIX Security Symposium (USENIX Security 25), 2025, pp. 763–782.
[206] R. Wails, R. Jansen, A. Johnson, M. Sherr, UPGen: The Unidentified Protocol Generator, 2025. https://github.com/unblockable/upgen.
[207] N. Niere, F. Lange, R. Merget, J. Somorovsky, Transport Layer Obscurity: Circumventing SNI Censorship on the TLS-Layer, in: 2025 IEEE Symposium on Security and Privacy (SP), IEEE, 2025, pp. 1344–1362.
[208] N. Niere, F. Lange, R. Merget, J. Somorovsky, Censor-Scanner, 2025. https://github.com/tls-attacker/Censor-Scanner/releases/tag/v1.0_sp2025.
[209] N. Tusing, J. Oakley, C. Shao, L. Yu, R. Brooks, Minecraft tunnels for covert communications, Entertain. Comput. 53 (2025) 100924.
[210] N. Tusing, J. Oakley, C. Shao, L. Yu, R. Brooks, Minecruft-PT, 2025. https://github.com/doudoulong/Minecruft-PT.
[211] S. Wendzel, T. Schmidbauer, S. Zillien, J. Keller, DYST (did you see that?): an amplified covert channel that points to previously seen data, IEEE Trans. Dependable Secure Comput. 22 (1) (2024) 614–631.
[212] S. Wendzel, T. Schmidbauer, S. Zillien, J. Keller, DYST (Did You See That?) – A Covert Channel Exploiting Recent Legitimate Traffic, 2024. https://github.com/NIoSaT/DYST.
[213] P. Vines, S. McKay, J. Jenter, S. Krishnaswamy, Raceboat, 2024. https://github.com/tst-race/raceboat/.
[214] P. Vines, Ten years gone: revisiting cloud storage transports to reduce censored user burdens, in: Free and Open Communications on the Internet (FOCI), 2024, pp. 34–41.
[215] P. Müller, N. Niere, F. Lange, J. Somorovsky, Turning Attacks into Advantages: Evading HTTP Censorship with HTTP Request Smuggling, Free and Open Communications on the Internet (FOCI) (2024) 2–53.

[216] P. Müller, N. Niere, F. Lange, J. Somorovsky, Turning Attacks into Advantages: Evading HTTP Censorship with HTTP Request Smuggling, 2024. https://github.com/UPB-SysSec/SmugglingCircumventionResults.
[217] S. Zillien, T. Schmidbauer, M. Kubek, J. Keller, S. Wendzel, Look what's there! utilizing the Internet's existing data for censorship circumvention with OPPRESSION, in: Proceedings of the 19th ACM Asia Conference on Computer and Communications Security, 2024, pp. 80–95.
[218] S. Zillien, T. Schmidbauer, M. Kubek, J. Keller, S. Wendzel, OPPRESSION, 2024. https://github.com/Stego-Punk-Lab/OPPRESSION.
[219] P.T.J. Kon, A. Gattani, D. Saharia, T. Cao, D. Barradas, A. Chen, M. Sherr, B.E. Ujcich, NetShuffle, 2024. https://github.com/patrickkon/NetShuffle.
[220] P.T.J. Kon, S. Kamali, J. Pei, D. Barradas, A. Chen, M. Sherr, M. Yung, SpotProxy: Rediscovering the Cloud for Censorship Circumvention, in: 33rd USENIX Security Symposium (USENIX Security 24), 2024, pp. 2653–2670.
[221] P.T.J. Kon, S. Kamali, J. Pei, D. Barradas, A. Chen, M. Sherr, M. Yung, SpotProxy, 2024. https://github.com/spotproxy-project/spotproxy.
[222] E. Chi, G. Wang, J.A. Halderman, E. Wustrow, J. Wampler, Just add water: Webassembly-based circumvention transports, Free and Open Communications on the Internet (1) (2024) 22–28.
[223] C. Bocovich, A. Breault, D.F. Serene, X. Wang, W.a.t.e.r.: Webassembly transport executables runtime, 2024. https://github.com/refraction-networking/water.
[224] A. Vilalonga, J.S. Resende, H. Domingos, Pub/Sub Rendezvous Protocol, 2024. https://github.com/AfonsoVilalonga/PubSub-Rendezvous.
[225] M. Pu, A. Wang, A. Chang, K. Quan, Y.W. Zhou, SQS Rendezvous Method, 2024. https://gitlab.torproject.org/tpo/anti-censorship/pluggable-transports/snowflake/-/merge_requests/214.
[226] W. Jia, M. Wang, L. Wang, P. Mittal, Quicstep: Circumventing quic-based censorship, arXiv preprint arXiv:2304.01073, 2023.
[227] W. Jia, M. Wang, L. Wang, P. Mittal, Quicstep, 2023. https://github.com/inspire-group/quicstep.
[228] R. Wails, R. Jansen, A. Johnson, M. Sherr, Proteus: Programmable Protocols for Censorship Circumvention, in: Workshop on Free and Open Communication on the Internet, 2023.
[229] R. Wails, R. Jansen, A. Johnson, M. Sherr, Proteus, 2023. https://github.com/unblockable/proteus.
[230] W. Jia, J. Eichenhofer, L. Wang, P. Mittal, Voiceover: Censorship-Circumventing Protocol Tunnels with Generative Modeling, Free and Open Communications on the Internet (FOCI) (2023) 67–80.
[231] W. Jia, J. Eichenhofer, L. Wang, P. Mittal, Voiceover: Censorship-Circumventing Protocol Tunnels with Generative Modeling, 2023. https://github.com/watsonjia/voiceover.
[232] J. Ding, K. Chen, Y. Wang, N. Zhao, W. Zhang, N. Yu, Discop: Provably secure steganography in practice based on "distribution copies", 2023. https://github.com/comydream/Discop.
[233] R. Wails, A. Stange, E. Troper, A. Caliskan, R. Dingledine, R. Jansen, M. Sherr, Learning to Behave: Improving Covert Channel Security with Behavior-Based Designs, in: Proceedings on Privacy Enhancing Technologies, 2022.
[234] G. Figueira, D. Barradas, N. Santos, Stegozoa: Enhancing WebRTC covert channels with video steganography for Internet censorship circumvention, in: Proceedings of the 2022 ACM on Asia Conference on Computer and Communications Security, 2022, pp. 1154–1167.
[235] G. Figueira, D. Barradas, N. Santos, Stegozoa, 2022. https://github.com/GabrielCFigueira/stegozoa-video.
[236] A. Devraj, L. Wang, J. Rexford, REDACT: refraction networking from the data center, ACM SIGCOMM Comput. Commun. Rev. 51 (4) (2021) 15–22.
[237] A.H. Lorimer, L. Tulloch, C. Bocovich, I. Goldberg, OUStralopithecus: Overt user simulation for censorship circumvention, in: Proceedings of the 20th Workshop on Workshop on Privacy in the Electronic Society, 2021, pp. 137–150.
[238] A.H. Lorimer, L. Tulloch, C. Bocovich, I. Goldberg, OUStral-library, 2021. https://gitlab.com/oustral/oustral.
[239] G. Kaptchuk, T.M. Jois, M. Green, A.D. Rubin, Meteor: Cryptographically secure steganography for realistic distributions, in: Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security, 2021, pp. 1529–1548.
[240] G. Kaptchuk, T.M. Jois, M. Green, A.D. Rubin, meteor, 2021. https://meteorfrom.space/.
[241] S. Satija, R. Chatterjee, BlindTLS: Circumventing TLS-based HTTPS censorship, in: Proceedings of the ACM SIGCOMM 2021 Workshop on Free and Open Communications on the Internet, 2021, pp. 43–49.
[242] M.B. Rosen, J. Parker, A.J. Malozemoff, Balboa: Bobbing and weaving around network censorship, in: 30th USENIX Security Symposium (USENIX Security, vol. 21, 2021, pp. 3399–3413.
[243] A.J. Malozemoff, C. Brent, J. Parker, M. Rosen, Shyamshankar, rocky: A rust library suite for channel obfuscation, 2021. https://github.com/GaloisInc/balboa.
[244] P.K. Sharma, D. Gosain, S. Chakravarty, Camoufler: Accessing the censored web by utilizing instant messaging channels, in: Proceedings of the 2021 ACM Asia Conference on Computer and Communications Security, 2021, pp. 147–161.
[245] M. Wei, Domain Shadowing: Leveraging Content Delivery Networks for Robust Blocking-Resistant Communications, in: 30th USENIX Security Symposium (USENIX Security, vol. 21, 2021, pp. 3327–3343.
[246] B. Birtel, C. Rossow, Slitheen++: Stealth TLS-based Decoy Routing, in: 10th USENIX Workshop on Free and Open Communications on the Internet (FOCI 20), 2020.
[247] B. Birtel, C. Rossow, Slitheen++, 2020. https://archive.org/details/Slitheen_plus_plus_source_code_20200901.
[248] M. Minaei, P. Moreno-Sanchez, A. Kate, Moneymorph: Censorship resistant rendezvous using permissionless cryptocurrencies, in: Proceedings on Privacy Enhancing Technologies 2020, 2020, pp. 3.
[249] M. Minaei, P. Moreno-Sanchez, A. Kate, Moneymorph: Censorship resistant rendezvous using permissionless cryptocurrencies, 2020. https://github.com/moneymorph.
[250] Y. Govil, L. Wang, J. Rexford, MIMIQ: Masking IPs with Migration in QUIC, in: 10th USENIX Workshop on Free and Open Communications on the Internet (FOCI 20), 2020.
[251] Y. Govil, L. Wang, J. Rexford, MIMIQ, 2020. https://github.com/ygovil1/mimiq.
[252] D. Barradas, N. Santos, Towards a scalable censorship-resistant overlay network based on WebRTC covert channels, in: Proceedings of the 1st International Workshop on Distributed Infrastructure for Common Good, 2020, pp. 37–42.
[253] D. Barradas, N. Santos, L. Rodrigues, V. Nunes, Poking a hole in the wall: Efficient censorship-resistant Internet communications by parasitizing on WebRTC, in: Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security, 2020, pp. 35–48.
[254] D. Barradas, N. Santos, L. Rodrigues, V. Nunes, Protozoa, 2020. https://github.com/dmbb/Protozoa.
[255] D. Fifield, Turbo Tunnel, a good way to design censorship circumvention protocols, in: 10th USENIX Workshop on Free and Open Communications on the Internet (FOCI 20), 2020.
[256] D. Fifield, Turbo Tunnel, 2020. https://www.bamsoftware.com/papers/turbotunnel/.
[257] P.K. Sharma, D. Gosain, H. Sagar, C. Kumar, A. Dogra, V. Naik, H.B. Acharya, S. Chakravarty, Siegebreaker: An SDN based practical decoy routing system, in: Proceedings on Privacy Enhancing Technologies, 2020.
[258] P.K. Sharma, D. Gosain, H. Sagar, C. Kumar, A. Dogra, V. Naik, H.B. Acharya, S. Chakravarty, SiegeBreaker, 2020. https://github.com/pi-yush/SiegeBreaker.
[259] M. Nasr, H. Zolfaghari, A. Houmansadr, A. Ghafari, MassBrowser: Unblocking the Censored Web for the Masses, by the Masses, in: Network and Distributed Systems Security (NDSS) Symposium 2020, 2020.
[260] M. Nasr, H. Zolfaghari, A. Houmansadr, A. Ghafari, MassBrowser, 2020. https://massbrowser.cs.umass.edu/.
[261] B. VanderSloot, S. Frolov, J. Wampler, S.C. Tan, I. Simpson, M. Kallitsis, J.A. Halderman, N. Borisov, E. Wustrow, Running refraction networking for real, in: Proceedings on Privacy Enhancing Technologies, 2020, pp. 321–335.
[262] S. Frolov, F. Douglas, W. Scott, A. McDonald, B. VanderSloot, R. Hynes, A. Kruger, M. Kallitsis, D.G. Robinson, S. Schultze, et al., TapDance, 2017. https://github.com/refraction-networking/tapdance.
[263] Z. Wang, S. Zhu, SymTCP: Eluding stateful deep packet inspection with automated discrepancy discovery, in: Network and Distributed System Security Symposium (NDSS), 2020.
[264] Z. Wang, S. Zhu, SymTCP: Eluding stateful deep packet inspection with automated discrepancy discovery, 2020. https://github.com/seclab-ucr/SymTCP.
[265] J. Oakley, L. Yu, X. Zhong, G.K. Venayagamoorthy, R. Brooks, Protocol proxy: An FTE-based covert channel, Comput. Secur. 92 (2020) 101777.
[266] S. Frolov, J. Wampler, S.C. Tan, J.A. Halderman, N. Borisov, E. Wustrow, Conjure: Summoning proxies from unused address space, in: Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security, 2019, pp. 2215–2229.
[267] S. Frolov, J. Wampler, S.C. Tan, J.A. Halderman, N. Borisov, E. Wustrow, Refraction Networking Client, 2019. https://github.com/refraction-networking/gotapdance/tree/dark-decoy.
[268] S. Frolov, E. Wustrow, The use of TLS in Censorship Circumvention, in: Network and Distributed Systems Security (NDSS) Symposium 2019, 2019, pp. 1–15.
[269] S. Frolov, E. Wustrow, uTLS, 2019. https://github.com/refraction-networking/utls.
[270] C. Bocovich, I. Goldberg, Secure asymmetry and deployability for decoy routing systems, in: Proceedings on Privacy Enhancing Technologies, vol. 3, 2018, pp. 43–62.
[271] V. Manfredi, P. Songkuntham, MultiFlow: Cross-Connection Decoy Routing using TLS 1.3 Session Resumption, in: 8th USENIX Workshop on Free and Open Communications on the Internet (FOCI 18), 2018.
[272] M. Nasr, H. Zolfaghari, A. Houmansadr, The waterfall of liberty: Decoy routing circumvention that resists routing attacks, in: Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, 2017, pp. 2037–2052.
[273] M. Nasr, H. Zolfaghari, A. Houmansadr, Waterfall, 2017. https://github.com/ha-D/waterfall.
[274] D. Barradas, N. Santos, L. Rodrigues, DeltaShaper: Enabling unobservable censorship-resistant TCP tunneling over videoconferencing streams, in: Proceedings on Privacy Enhancing Technologies, 2017.
[275] D. Barradas, N. Santos, L. Rodrigues, DeltaShaper, 2017. https://github.com/dmbb/DeltaShaper.
[276] V. Heydari, S.-I. Kim, S.-M. Yoo, Scalable anti-censorship framework using moving target defense for web servers, IEEE Trans. Inf. Forensics Secur. 12 (5) (2017) 1113–1124.
[277] S. Frolov, F. Douglas, W. Scott, A. McDonald, B. VanderSloot, R. Hynes, A. Kruger, M. Kallitsis, D.G. Robinson, S. Schultze, et al., An ISP-Scale Deployment of TapDance, in: 7th USENIX Workshop on Free and Open Communications on the Internet (FOCI 17), 2017.
[278] Z. Li, S. Herwig, D. Levin, DeTor: Provably Avoiding Geographic Regions in Tor, 2017. https://detor.cs.umd.edu/.

[279] C. Bocovich, I. Goldberg, Slitheen: Perfectly imitated decoy routing through traffic replacement, in: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, 2016, pp. 1702–1714.
[280] C. Bocovich, I. Goldberg, Slitheen, 2016. https://gitlab.com/slitheen.
[281] R. McPherson, A. Houmansadr, V. Shmatikov, CovertCast: Using Live Streaming to Evade Internet Censorship, in: Proceedings on Privacy Enhancing Technologies, 2016.
[282] R. McPherson, A. Houmansadr, V. Shmatikov, CovertCast, 2016. https://github.com/rfmcpherson/CovertCast.
[283] V. Heydari, S.-I. Kim, S.-M. Yoo, Anti-censorship framework using mobile IPv6 based moving target defense, in: Proceedings of the 11th Annual Cyber and Information Security Research Conference, 2016, pp. 1–8.
[284] H. Zolfaghari, A. Houmansadr, Practical censorship evasion leveraging content delivery networks, in: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, 2016, pp. 1715–1726.
[285] J. Holowczak, A. Houmansadr, CacheBrowser, 2015. https://github.com/CacheBrowser.
[286] S. Li, N. Hopper, Mailet: Instant social networking under censorship, Proc. Priv. Enhancing Technol. (2) (2016) 175–192.
[287] S. Li, N. Hopper, Mailet: Instant social networking under censorship, 2016. https://github.com/magicle/Mailet.
[288] M. Nasr, A. Houmansadr, Game of decoys: Optimal decoy routing through game theory, in: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, 2016, pp. 1727–1738.
[289] Q.U.A.D. Akbar, M. Flores, A. Kuzmanovic, DNS-sly: Avoiding censorship through network complexity, in: 6th USENIX Workshop on Free and Open Communications on the Internet, 2016.
[290] B. Hahn, R. Nithyanand, P. Gill, R. Johnson, Games without frontiers: Investigating video games as a covert channel, in: 2016 IEEE European Symposium on Security and Privacy (EuroS&P), IEEE, 2016, pp. 63–77.
[291] B. Hahn, R. Nithyanand, P. Gill, R. Johnson, Castle-Covert-Channel, 2016. https://github.com/bridgar/Castle-Covert-Channel.
[292] F. Douglas, W. Pan, M. Caesar, et al., Salmon: Robust proxy distribution for censorship circumvention, in: Proceedings on Privacy Enhancing Technologies, 2016.
[293] F. Douglas, W. Pan, M. Caesar, et al., Salmon Project, 2016. https://github.com/SalmonProject.
[294] K. Kohls, T. Holz, D. Kolossa, C. Pöpper, SkypeLine: Robust hidden data transmission for VoIP, in: Proceedings of the 11th ACM on Asia Conference on Computer and Communications Security, 2016, pp. 877–888.
[295] P. Vines, T. Kohno, Rook, 2015. https://github.com/plvines/rook.
[296] D. Levin, Y. Lee, L. Valenta, Z. Li, V. Lai, C. Lumezanu, N. Spring, B. Bhattacharjee, Alibi routing, 2015. https://alibi.cs.umd.edu/.
[297] J. Holowczak, A. Houmansadr, Cachebrowser: Bypassing Chinese Censorship without Proxies using Cached Content, in: Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, 2015, pp. 70–83.
[298] D. Ellard, C. Jones, V. Manfredi, W.T. Strayer, B. Thapa, M. Van Welie, A. Jackson, Rebound: Decoy routing on asymmetric routes via error messages, in: 2015 IEEE 40th Conference on Local Computer Networks (LCN), IEEE, 2015, pp. 91–99.
[299] D. Ellard, C. Jones, V. Manfredi, W.T. Strayer, B. Thapa, M. Van Welie, A. Jackson, Curveball Snapshots, 2015. https://github.com/DanielEllard/curveball-snapshots.
[300] K.P. Dyer, S.E. Coull, T. Shrimpton, Marionette: A programmable network traffic obfuscation system, in: 24th USENIX Security Symposium (USENIX Security 15), 2015, pp. 367–382.
[301] K.P. Dyer, S.E. Coull, T. Shrimpton, marionette, 2015. https://github.com/marionette-tg/marionette.
[302] E. Wustrow, S. Wolchok, I. Goldberg, J.A. Halderman, Telex: Anticensorship in the network infrastructure, in: 20th USENIX Security Symposium (USENIX Security 11), 2011.
[303] J. Karlin, D. Ellard, A.W. Jackson, C.E. Jones, G. Lauer, D.P. Mankins, W.T. Strayer, Decoy routing: Toward unblockable Internet communication, in: USENIX workshop on free and open communications on the Internet (FOCI 11), 2011.
[304] M. Schuchard, J. Geddes, C. Thompson, N. Hopper, Routing around decoys, in: Proceedings of the 2012 ACM conference on Computer and communications security, 2012, pp. 85–96.
[305] D. Gosain, A. Agarwal, S. Chakravarty, H.B. Acharya, The devil's in the details: Placing decoy routers in the Internet, in: Proceedings of the 33rd Annual Computer Security Applications Conference, 2017, pp. 577–589.
[306] Cloudflare, What is a content delivery network (CDN)? 2025. https://www.cloudflare.com/learning/cdn/what-is-a-cdn/.
[307] Kaspersky, Domain shadowing, 2024. https://encyclopedia.kaspersky.com/glossary/domain-shadowing/.
[308] D. Gale, L.S. Shapley, College admissions and the stability of marriage, The American mathematical monthly 69 (1) (1962) 9–15.
[309] P. S. Université, PlanetLab, 2025. https://planetlab.io/.
[310] B. Schneier, Applied Cryptography, Second Edition: Protocols, Algorithms, and Source Code in C, second ed, John Wiley & Sons, Inc, New York, 1996, cloth edition.
[311] L. Wang, K.P. Dyer, A. Akella, T. Ristenpart, T. Shrimpton, Seeing through network-protocol obfuscation, in: Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, 2015, pp. 57–69.
[312] S. Fan, J. Sippe, S. San, J. Sheffey, D. Fifield, A. Houmansadr, E. Wedwards, E. Wustrow, Wallbleed: a memory disclosure vulnerability in the Great Firewall of China, in: Network and Distributed System Security, The Internet Society, 2025, pp. 1–20.
[313] N. Demir, M. Große-Kampmann, T. Urban, C. Wressnegger, T. Holz, N. Pohlmann, Reproducibility and replicability of web measurement studies, in: Proceedings of the ACM Web Conference 2022, 2022, pp. 533–544.
[314] F. Hantke, P. Snyder, H. Haddadi, B. Stock, arXiv preprint arXiv:2501.15911, 2025.
[315] A. IMC, Call For Papers, in: ACM IMC 2025, 2025, https://conferences.sigcomm.org/imc/2025/cfp/#replicability-track.
[316] F. House, Freedom on the Net 2025: Cuba: Key Developments, 2025, https://freedomhouse.org/country/cuba/freedom-net/2025 (1 June 2024).
[317] A. Mare, State-ordered internet shutdowns and digital authoritarianism in Zimbabwe, Int. J. Commun. 14 (2020) 20.
[318] E. Marchant, N. Stremlau, A spectrum of shutdowns: Reframing internet shutdowns from Africa, Int. J. Commun. 14 (2020).
[319] E. Marchant, N. Stremlau, The changing landscape of internet shutdowns in Africa—Introduction, Int. J. Commun. 14 (2020).
[320] S.Y. Wurah, N.B. Anya, T.A. Benson, A Call to Arms: Motivating An Internet Measurements Observatory for Africa, in: Proceedings of the 24th ACM Workshop on Hot Topics in Networks, 2025, pp. 104–113.
[321] F. Holzbauer, S. Strobl, J. Ullrich, Tracking Internet Disruptions in Ukraine: Insights from Three Years of Active Full Block Scans, in: Proceedings of the 2025 ACM Internet Measurement Conference, 2025, pp. 474–492.
[322] F. Cristiano, The blurring politics of cyber conflict: A critical study of the digital in Palestine and beyond, Lund Political Studies 206 (2022).
[323] W. Magdy, H. Mubarak, J. Salminen, Who should set the standards? Analysing censored Arabic content on Facebook during the Palestine-Israel conflict, in: Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems, 2025, pp. 1–16.