


Journal Pre-proof

A Survey of Internet Censorship and its Measurement: Methodology, Trends, and Challenges

Steffen Wendzel, Simon Volpert, Sebastian Zillien, Julia Lenz, Philip Rünz, Luca Caviglione

PII: S0167-4048(25)00421-3
DOI: <https://doi.org/10.1016/j.cose.2025.104732>
Reference: COSE 104732



To appear in: *Computers & Security*

Received date: 8 May 2025
Revised date: 5 September 2025
Accepted date: 21 October 2025

Please cite this article as: Steffen Wendzel, Simon Volpert, Sebastian Zillien, Julia Lenz, Philip Rünz, Luca Caviglione, A Survey of Internet Censorship and its Measurement: Methodology, Trends, and Challenges, *Computers & Security* (2025), doi: <https://doi.org/10.1016/j.cose.2025.104732>

This is a PDF file of an article that has undergone enhancements after acceptance, such as the addition of a cover page and metadata, and formatting for readability, but it is not yet the definitive version of record. This version will undergo additional copyediting, typesetting and review before it is published in its final form, but we are providing this version to give early visibility of the article. Please note that, during the production process, errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

© 2025 The Author(s). Published by Elsevier Ltd.
This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>)

A Survey of Internet Censorship and its Measurement: Methodology, Trends, and Challenges

Steffen Wendzel^a, Simon Volpert^a, Sebastian Zillien^a, Julia Lenz^a, Philip Rünz^b, Luca Caviglione^c

^a*Institute of Information Resource Management (IRM), Dep. Engineering, Computer Science and Psychology, Ulm University, Germany*

^b*Faculty of Mathematics & Computer Science, University of Hagen, Germany*

^c*National Research Council of Italy, Genoa, Italy*

Abstract

Internet censorship limits the access of nodes residing within a specific network environment to the public Internet, and vice versa. During the last decade, techniques for conducting Internet censorship have been developed further. Consequently, methodology for *measuring* Internet censorship had been improved as well.

In this paper, we firstly provide a survey of network-level Internet censorship techniques. Secondly, we survey censorship measurement methodology. We further cover the censorship of circumvention tools and its measurement, as well as available datasets. In cases where it is beneficial, we bridge the terminology and taxonomy of Internet censorship with related domains, namely traffic obfuscation and information hiding. We further extend the technical perspective with recent trends and challenges, including human aspects of Internet censorship.

Keywords: Internet Measurement, IPv6, TCP, UDP, QUIC, DNS, HTTPS, TLS, SNI, Tor, BGP, VPN, GFW, Censorship, Human Aspects of Information Security

1. Introduction

The Internet enables the interconnectivity of numerous users and the interaction with millions of websites and online services. Sometimes, content or user-interaction is considered inappropriate or undesired by entities controlling fractions of the Internet. Such entities can be state governments, regional governments and any form of corporate and non-governmental organizations. The desire to control – or: censor – content of the Internet for its users has been studied by a plethora of researchers. Internet censorship of countries has already been discussed in the mid-1990's [11] and early forms of Internet censorship have also been used in an *idealistic vision of self-governance* by the Internet community [136]. Since then, due to the increasing number of Internet users, social media platforms and the ever-growing amount of online content, Internet censorship became a widespread and persistent phenomenon [72].

While there are already some surveys on Internet censorship, none of these studies up-to-date techniques of both, Internet censorship *and* its measurement. Further, no previous work draws clear links between Internet censorship and the related domains of network traffic obfusca-

tion and network information hiding through covert channels and steganography. Additionally, previous surveys did not cover related human aspects.

To close this gap, we provide a survey on network-level Internet censorship techniques, their measurement and selected related topics. While we cover some human aspects, surveying political and juridical aspects of censorship (e.g., content removal by authorities) is outside the scope of our paper. We then only provide links to such topics when it aids the understanding.

The **key contributions of this article** are the following:

1. We perform survey on *Internet censorship methodology*, i.e., how it is technically realized on the network level.
2. We also survey network-level censorship *measurement*.
3. We discuss the censorship techniques applied to circumvention tools and the measurement of the related censorship techniques.
4. We survey measurement *datasets* and show how they overlap and where they have singularities.
5. Whenever reasonable, we draw links between Internet censorship and related domains, such as traffic obfuscation.
6. We summarize trends and challenges of Internet censorship and its measurement, including both, technical, societal and human aspects.

Email addresses: steffen.wendzel@uni-ulm.de (Steffen Wendzel), simon.volpert@uni-ulm.de (Simon Volpert), sebastian.zillien@uni-ulm.de (Sebastian Zillien), julia.lenz@uni-ulm.de (Julia Lenz), philip.ruenz@studium.fernuni-hagen.de (Philip Rünz), luca.caviglione@cnr.it (Luca Caviglione)

The rest of the paper is structured as follows. Sect. 2 covers the related work. Sect. 3 describes our literature selection procedure while Sect. 4 introduces fundamentals and a taxonomy of censor capabilities. Next, Sect. 5 surveys censorship techniques, whereas Sect. 6 reviews mechanisms to detect censorship. We cover the censorship of circumvention tools and the measurement for these tools' censorship in Sect. 7. Measurement platforms and datasets are summarized in Sect. 8. Sect. 9 discusses trends and challenges and Sect. 10 draws conclusions.

2. Related Work

Surveys on Internet censorship have been published by several authors. We will first summarize the key characteristics of these surveys and then show how our paper improves over these existing surveys. A fraction of the available surveys discusses (mostly) political aspects of Internet censorship (e.g., [216, 23]) but we decided to exclude such works due to the technical focus of our paper.

Tab. 1 provides an overview of the existing surveys in the field, highlighting their major focus. We compare the Internet protocols covered by the various reviews and censorship detection techniques in Tab. 2.

In a 2011-survey, Subramanian [182] summarizes heterogeneous aspects of Internet censorship for selected countries, including circumvention methods. Aceto and Pescapé [2] provide a comprehensive survey on Internet censorship detection but also cover censorship methodologies. Dated 2015, the survey covers a plethora of fundamental methods for censorship detection but lacks developments of the last ten years. Moreover, the authors neither discuss information hiding and traffic obfuscation methods nor do they cover newer protocols (e.g., IPv6 and QUIC) or human aspects of Internet censorship. Tschantz *et al.* [190] survey censorship *circumvention* approaches, especially their related evaluation methodology. They further discuss circumvention approaches within the context of real-world censor's known capabilities of 2016 and conclude that there is a disconnect between the goals of sophisticated circumvention methods and the censor's methods. They present three key reasons for this mismatch, i.e., (i) that censors focus on the discovery and setup procedure of circumvention channels while circumvention research mostly focuses on the channel itself; (ii) that censors apply cheap passive censorship approaches instead of complex ones; and (iii) that censors try not to risk false blocking while research covers several attacks with collateral damage [190]. In their textbook, Mazurczyk *et al.* [128] discuss network information hiding methods, including network steganography and traffic obfuscation. Regarding the countermeasures, the work mostly focuses on the detection, elimination, and prevention of covert channels on different levels of the TCP/IP stack, including specific covert channel/circumvention tools. Khattak *et al.* [103] survey censorship-resistant communication tools both from practical and theoretical viewpoints. A 2017-survey by Zittrain *et al.* [213]

conducts a comparison of censorship on different nations and also investigates the impact of HTTPS on Internet censorship. While not a survey *per se*, Niaki *et al.* provide an overview of several censorship measurement methods employed by ICLAB [140], featuring approaches for prominent communication protocols. Ververis *et al.* conduct a cross-country comparison of Internet censorship [192]. Their survey focuses on France, Turkey, and Iran. Karunanayake *et al.* [101] provide a comprehensive survey of de-anonymization attacks on Tor. Their work features a well-grounded taxonomy of Tor attacks. Master and Garman [124] conduct a cross-sectional study on Internet censorship of 70 countries. They also summarize the applied censorship methods and found that censors still apply filtering methods that are easy to bypass and perform total Internet shutdowns. Hall *et al.* [82] provide a survey-styled RFC covering several network-level aspects of Internet censorship methods. While obfuscation methods are briefly mentioned, information hiding techniques, measurement techniques, and human aspects are largely neglected. Finally, Wrana *et al.* study the potential for censorship in future Internet architectures, including NDN, SCION, XIA, and NEBULA.

In comparison to the above-mentioned surveys, we do not only provide the most up-to-date coverage of censorship and measurement methodology (including several 2025-publications and novel trends, such as regional censorship or detection of vulnerabilities in censorship systems) that features the broadest set of network protocols. We also draw links to related disciplines, such as information hiding. Finally, to the best of the authors knowledge, our survey is the only one that explicitly covers human aspects linked to Internet censorship.

3. Literature Selection

Our literature selection process is sketched in Fig. 1. First, we conducted a query on the dblp computer science literature repository for the keywords **censor*** AND (**network** OR **internet** OR **block*** OR **firewall** OR **circumvent*** OR **measure***). We considered papers that had been published in the 2010-2024 range. Additionally, we queried Google Scholar and Web of Science (WoS) with the term **Internet censorship** (in both cases, we sorted the listing of results "by relevance", excluded papers dated earlier than 2010, and considered the first 250 results). All queries were conducted in October 2024 and redundancies were removed. Note that we included pre-prints/technical reports. Next, we manually monitored high-rank conferences and journals as well as specialized events, such as FOCI, for freshly appearing papers, which led to the inclusion of several late-2024- and 2025-papers. All found papers were filtered, i.e., we excluded papers on political aspects of censorship, papers for which other attributes (such as author names) led to their inclusion, papers on censorship circumvention approaches (without discussing

Table 1: Comparison of Existing Surveys’ Major Focus

Author(s)	Ref.	Year	Published	Main Focus	Int. Cens.	Inf. Hid.	Traf. Obs.
Subramanian	[182]	2011	CIIMA	General Internet censorship	✓	-	-
Aceto and Pescapé	[2]	2015	Comp. Netw.	Censorship detection	✓	-	-
Tschantz <i>et al.</i>	[190]	2016	IEEE S&P	Evaluation of circumvention methods	✓	✓	✓
Mazurczyk <i>et al.</i>	[128]	2016	textbook	Network information hiding	-	✓	✓
Khattak <i>et al.</i>	[103]	2016	PoPETS	Censorship circumvention systems	✓	✓	✓
Zittrain <i>et al.</i>	[213]	2017	research report	Country-based comparison; HTTPS	✓	-	-
Niaki <i>et al.</i>	[140]	2020	IEEE S&P	Censorship measurement	✓	-	-
Ververis <i>et al.</i>	[192]	2020	Policy & Internet	Country-based comparison of censorship	✓	-	-
Karunanayake <i>et al.</i>	[101]	2021	COMST	Tor de-anonymization attacks	✓	-	-
Master and Garman	[124]	2023	FOCI	Current censorship in 70 countries	✓	-	-
Hall <i>et al.</i>	[82]	2023	IRTF RFC	Internet censorship methods	✓	-	-
Wrana <i>et al.</i>	[205]	2025	PETS	Future Internet architectures	✓	(✓)	-
This Paper	-	2025	-	Censorship and its measurement	✓	✓	✓

Table 2: Comparison of Network Protocol-based Techniques Covered by Existing Surveys. (Checkmarks in brackets indicate that the topic is partially addressed but not fully surveyed. This case appeared mostly when papers had a different focus, such as performing a cross-country analysis.)

General Aspects			Censorship Meth./Countermeasures											Measurement Meth.										
Author(s)	Ref.	Coverage of last five years	Taxonomy	IPv4	IPv6	BGP	TCP, UDP, QUIC	DNS	HTTP(S)	TLS (or: SSL)	Circumvention Tools / VPNs	Other Appl. Protocols	Active Probing	Taxonomy	IPv4	IPv6	BGP	TCP, UDP, QUIC	DNS	HTTP(S)	TLS (or: SSL)	Circumvention Tools / VPNs	Other Appl. Protocols	Active Probing
Subramanian	[182]	-	✓	✓	-	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Aceto and Pescapé	[2]	-	✓	✓	-	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Tschantz <i>et al.</i>	[190]	-	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Mazurczyk <i>et al.</i>	[128]	-	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Khattak <i>et al.</i>	[103]	-	✓	✓	-	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Zittrain <i>et al.</i>	[213]	-	✓	✓	-	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Niaki <i>et al.</i>	[140]	-	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Ververis <i>et al.</i>	[192]	-	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Karunanayake <i>et al.</i>	[101]	1y	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Master and Garman	[124]	3y	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Hall <i>et al.</i>	[82]	3y	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Wrana <i>et al.</i>	[205]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
This Paper	-	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

relevant measurement aspects) and papers on entirely different topics that matched our keywords. In addition, we performed a snowballing procedure with all survey papers that we found as well as with papers cited at least 50 times according to WoS to detect additional relevant works that could be included. Moreover, we manually scanned Google Scholar for relevant yet dated papers that appeared before 2010 so that we were able to put current techniques into historic context. All papers were read, and a team discussion was held to decide on the exclusion of border-case papers. Thankfully, the Reviewers pointed us towards a small number of papers that had not been found by our methodology, such as RFC 9505 [82]. These papers had been included in the present version of the work.

4. Fundamentals

During the years, Internet censorship has developed its own jargon and also started a vivid research area overlapping across several disciplines (e.g., privacy, network security and traffic engineering). In the following, we present the basic terminology as well as the capabilities of a censor organized within a suitable taxonomy.

4.1. Basic Terminology

Aceto and Pescapé refer to **Internet censorship** as the *intentional impairing or blocking of access to online resources and services* [2]. They further define **censorship circumvention** as the *process of nullifying the censoring action, i.e., accessing the unmodified target – or an equivalent copy – despite the presence of a censoring system*. In this context, the **target** is an *online resource or service*.




Figure 1: Paper selection methodology

The target can refer to nodes/endpoints, sites/URLs, protocols and other information that is required for a user to access it, while the altered experience of end-users when accessing a target is called a **symptom** [2].

Note that while Internet censorship, in a broader sense, also covers governmental/authority actions that moderate or remove content, we focus on pure network-level censorship and thus exclude political aspects of Internet censorship. However, we highlight relevant human aspects at the end of the paper.

Censorship can also be described by Simmons' *Prisoners' Problem* [181]. In this setting, Alice and Bob are prisoners, located in isolated cells. The only way they can exchange messages is through the **warden**. The warden can read and manipulate/drop these messages. The goal of Alice and Bob is to communicate by handing slightly manipulated messages to the warden. The manipulations represent secret information. The warden delivers these benign-looking messages while not recognizing their embedded secret message. In a censorship setting, Alice would aim to access a target (Bob) while the censor (the warden) influences that access. The term *warden* is widely used for censoring actions but also for adversaries and countermeasures that target censorship circumvention.

Hall *et al.* split the censorship process into three stages: *prescription*, *identification*, and *interference* [82] that correspond to the phases in which censors determine the targets they want to block, how they detect if users try accessing such targets, and how they handle access attempts to these targets, respectively. Following Aceto *et al.*, censorship can also be split into client-, server-, network-side and self-censorship [2]. **Client-side** censorship refers to




Figure 2: A censor's fundamental influence on communication between nodes

censorship actions performed on the network user's client, e.g., by means of a keyword filter or a personal firewall that blocks access to targets. Instead of such *additional* software components, censors can also offer **replacements for popular tools**, such as an own instant messaging or video conferencing tool with censorship functionalities already built-in. In contrast, **server-side** censorship is integrated into a (target's) server or applied to it. A technical example is a server-internal source address filter that prohibit access to selected resources for national users. **Network-side** censorship operates as part of the routing infrastructure. For instance, a censor could drop network packets that aim to reach undesired destinations. Similarly, undesired flows can be slowed down or redirected to monitoring networks. Finally, one form of censorship is **self-censorship**. In such a case, a user *self-restrict[s] his/her possibilities of the expression due to fear of punishment, retaliation, or other negative consequences* [2]. It is important to note, however, that self-censorship should be understood as a behavioral reaction to existing (perceived) censorship and surveillance, not as an intentional strategy imposed by censors. At the same time, it reinforces and extends the reach of censorship by amplifying its effects.

4.2. General Capabilities of a Censor

The capabilities of censors reportedly vary [103]. In general, a censor can both influence the communication within its own network infrastructure (Fig. 2(a)) or the communication to and from external nodes with internal ones (Fig. 2(b)). Note that the device that separates the censored from the public network is called the **network perimeter**, and its filtering actions **perimeter filtering** [199].

Khattak *et al.* describe a censor in an abstract model that (i) classifies traffic to detect flows that are relevant for blocking attempts; (ii) contains a cost function that considers collateral damage to a potential censorship action; and (iii) a decision function for determining the censorship action to be taken under consideration of outputs from the classifier and the cost function [103].

In general, a censor can perform four types of fundamental censoring actions [2, 128, 82]:




Figure 3: High-level censorship system taxonomy

1. **eliminate** the target’s content (e.g., taking down a particular website), which is usually caused by an administrative order.
2. **prevent access** to a non-eliminated target (e.g., blocking access to a host or the resolution of related DNS records, blocking VPN protocols and circumvention tools used to bypass filters, deactivating Internet access in a whole region etc.).
3. **limit** the communication with a target (e.g., slowing down a live stream of a protest), this is sometimes called **soft blocking**.
4. **monitor (fingerprint) behavior** and identify end-users by passive or active means (e.g., plain eavesdropping, redirecting traffic through analytical networks, or injecting a Trojan that leaks monitored data).

Taxonomy of censorship methods. Censorship actions can be categorized in a more fine-grained manner. For this purpose, we reviewed and merged both, literature on censorship methodology [2, 93] and on network information hiding [127, 128]. Such a merge turned out to be effective, especially due to the overlap between the two research areas. The resulting taxonomy is shown in Fig. 3.

First of all, we categorize censorship systems based on their (known) **censoring resources**. The fact that resources are a key issue for censors (and other complex filtering systems) was mentioned by several papers, e.g., [2, 127, 206]. Censorship systems can either be small (e.g., a SME, a single NGO, a local government) or large/state-level (including regional/national governments). This aspect is driven by the proposal of Houmansadr *et al.* to distinguish between local, state-level oblivious and state-level omniscient censors [93], which was also adapted to the case of network information hiding in [127]. Note that state-level oblivious and omniscient censors differ in terms

of their resources (the latter can *aggregate data collected at different network locations and store all intercepted traffic for offline, computationally expensive analysis* [93]). However, for our high-level view, it is enough to distinguish between the two cases mentioned above. Further, these two types of censor resource categories can employ **single-stage or a multi-stage censoring** (where a first-stage performs quick and superficial flow analysis and follow-up stages perform analyses in depth) [82, 199, 93].

While Aceto and Pescapé already include **localization** aspects, such as symptoms and triggers, the relevant generic aspects are the surveillance location (in the sense of technical monitoring) and the censoring location of the censorship system. Both can be either **in-line/on-path** (some authors differentiate between in-path and on-path, depending on the possibility of the censor to *modify* packets or passively monitoring them [197]) or **off-line/off-path** [82],[140, 99]. However, we extended the categorization with ideas of Mazurczyk *et al.* [127], where they differentiate between **centralized and distributed** systems. They further differentiate between localhost/intermediate node and network systems, which we merged with Aceto and Pescapé’s differentiation of client-side, server-side and network-side censorship, so that we finally have the options **client-local, intermediate node(s)** and **server-side**.

Regarding the **behavior** of the censorship system, most works (e.g., [2, 82, 197]) differentiate the **state-keeping** functionality (either **stateless** or **stateful**). However, only few works explicitly considered different “**consistencies**” (or: “**dynamics**”) [127, 197]. In this category, censorship can either be **static** (i.e., the behavior of the censorship system remains (mostly) unchanged over time) or **dynamic** (the behavior of the censorship system changes over time, e.g., to increase uncertainty of circumventing tool users or to rapidly adjust to new circumvention trends).

Finally, following Niaki *et al.* [140], censorship systems can operate in an **overt** way, i.e., they inform the user that censorship is happening, for instance through a blocking web page, or **covertly**, i.e., causing misleading errors so that the user does not recognize that censorship happens. Accordingly, we call this category **user-recognizability**.

Censorship Costs. Censorship is linked to certain (potential) costs, of which we identified three *major* categories:

1. Blocking and limitation techniques are often based on heuristics. Such heuristics are never “perfect” and might cause undesired side effects, e.g., due to imperfect regular expressions used in a filter. Moreover, there is an imperfect target collection procedure (e.g., incomplete or outdated list of IP address /websites to block). The two costs (side effects) are (i) preventing Internet access to targets not meant to be blocked (**overblocking** due to *false-positives*), and (ii) not covering all desired targets (**underblocking** due to *false-negatives*) [124, 132, 190, 199]. A

special form of overblocking called *censorship leakage* occurs when international traffic transits through a country but is unintentionally influenced by that country's censorship [117].

2. Monitoring/blocking flows and limiting throughput results in **resource costs** (equipment, staff, investments and operational costs) [204, 199, 190]. For this reason, a censor might not be able to run heuristics and blocking attempts on every packet or flow that passes through its network infrastructure [93, 138, 70, 132]. Thus, there are often at least two-stages of blocking: a non-detailed level with quick heuristics that are able to check many flows/packets. More interesting flows are then analyzed in-depth by more costly means, e.g., caching flows and conduct expensive post-analysis [93, 199]. Blocking *all* traffic (or a large fraction), e.g., during the *Arab Spring*, is often not considered attractive for the censor although was practically applied [132, 199]. A study performed in 2023 has shown that total Internet shutdown is still applied in several countries [124].
3. Censorship measures **decrease the reputation** of an entity acting as a censor, as well as the trust in it.

5. Network-level Censorship Techniques

A censor can select censorship techniques that target the users, hardware, software, or the network. For instance, country-level “real name” policies [2] can foster self-censorship. Hardware delivered to end-users could be manipulated or replaced before it reaches the customer. Moreover, a software manipulation can be realized by deploying a backdoor in its code.

In contrast to these methods, we focus on *network-level* censorship targeting the Internet, transport and application layers. In the network context, a censor could interrupt, redirect, alter, or record a node's transmission to specific destinations, e.g., through the physical manipulation of cables [124]. Most methods focus on the network level protocols above the physical and data link layers. This is why we survey censorship techniques residing on the Internet layer or above. After introducing the various censorship techniques, we also explain their limitations.

5.1. IPv4/IPv6-based Methodology

One of the most fundamental methods for Internet censorship is **discarding IPv4/IPv6 packets** destined to an undesired target [2, 103]. The target can reside outside the censor's network (such attempts have been made early on, see, e.g., [11]). Moreover, the censor can prevent communication of selected nodes *within* its network (e.g., nodes of a group of journalists and their chat service node). The censor can also **re-route** such packets through an analysis network for further **inspection**. This can be achieved through altered routing tables on a router or by sending

ICMP redirect messages (ICMPv4 type 5 / ICMPv6 type 137). However, ICMP redirects are often ignored by hosts who receive these messages due to security reasons.

A censor can also **introduce a disruptive ICMPv4 or ICMPv6 error response**, e.g., ICMPv4 type 3 or ICMPv6 type 1 (*destination unreachable*) with desired code options (*network / host / port unreachable*, *destination network unknown*, or *communication administratively prohibited*), or ICMPv4 type 11 / ICMPv6 type 3 (*time exceeded*), even if the packet would actually reach its destination [165, 2]. Finally, the censor can **decrease the Quality of Service (QoS) for selected flows** between sources and targets, which can be done by increasing packet loss, delays, or jitter as well as by reducing the bandwidth [2, 165, 103, 82].

5.2. Routing Protocol-based Methodology

Routing-based censorship techniques focus almost exclusively on the *Border Gateway Protocol* (BGP), which is the major routing protocol in the modern Internet. BGP routers exchange information about how they reach destination networks, taking care of the routing at large, i.e., between *Autonomous Systems* (AS) [42]. Therefore, BGP routers propagate *prefixes* (network identifier with a network mask) jointly with distance information.

Censorship methods for BGP have been summarized already by Aceto *et al.* [2]. The first method is to let a country's ASes **disappear** from the Internet by preventing the propagation of BGP updates, with users experiencing ICMP *network unreachable* errors. A second method is *prefix hijacking* [15, 103, 42, 2]. Traffic that should be directed towards an AS *B* is redirected to an AS *C* instead—through false propagation of a prefix not owned by the censor [15, 42]. A censor would need to propagate a short path for the hijacked prefix, as this increases the chances for neighbor routers to consider the route to *C* [42]. This allows for **blackhole attacks** so that a censor renders specific destinations unreachable for end-users [2]. Alternatively, a censor might transparently **redirect** traffic through its own AS *C* for **inspection**, and finally forwards the traffic to *B* so that users barely notice any alterations or delay [15, 42]. Users could experience *time exceeded* errors if redirected routing paths are too long [2]. While such blackholing and interception attacks can appear between ASes, a regional censor might also apply censoring methods on regional or even local routing environments, potentially even those running the *Open Shortest Path First* (OSPF) protocol.

5.3. UDP-, TCP- and QUIC-based Methodology

UDP- and TCP-based censorship techniques are often applied in the form of **port filtering**, and usually combined with specific IP addresses, forming tuples, such as (client IP, client port, server IP, server port) that are (temporarily) blocked [124]. As in the case of other protocols, **eavesdropping and consecutive traffic analysis** can also be conducted on the transport layer, e.g.,

to fingerprint the type of end-user device or the involved application-layer tool through packet sizes or other meta-data. Alternatively, undesired flows can also be **throttled**, e.g., by dropping selected UDP datagrams or delaying TCP acknowledgment segments on a router that is located on the path between end-user and target. Further, **hijacking** of transport-layer flows can happen at the start of a flow/connection or later, e.g., during or right after a TCP handshake was completed or in the middle of a download when some keywords were detected in the payload. Undesired datagrams and segments can be discarded and connections terminated (e.g., injecting spoofed TCP segments with active RST or FIN flags to terminate/close connections) [2, 165, 103]. Elmenhorst *et al.* report **blackholing** against QUIC connections that interrupt the handshake [60]. The resulting symptom for an end-user depends on the time of tampering. For instance, dropping packets during the TCP connection establishment/QUIC handshake phase lets the service appear unavailable (e.g., causing a timeout) [60] but dropping packets in the middle of a file download would let the connection appear unstable.

5.4. DNS-based Methodology

Censors can modify DNS servers as well as DNS entries (called *resource records*) in multiple ways. The most relevant types of resource records for censors are typically **A** (provides the IPv4 address of a hostname) and **AAAA** (provides the IPv6 address of a hostname). According to Aceto *et al.* and Khattak *et al.*, resource records of a censor-controlled DNS server can be influenced in several ways: (i) **resource record removal** (technical deletion from the DNS zone database) or **pretended in-existence** (sending an NXDOMAIN response), even if the resource record could be fetched from an external authoritative DNS server (that is explicitly not contacted by the censor's DNS server), (ii) **response alteration**: resource record's resolution can be modified so that it refers to an IP address that provides a "block page" or an "error page" (from an end-user perspective, only the content of the provided website differs), (iii) **returning a "failing IP"**, e.g., a non-routable in-house IP address, and (iv) **returning a network-level surveillance system's IP address** to further investigate traffic of clients [2, 103]. Such a surveillance system can also act as a transparent proxy that forwards all traffic to a client's desired destination (and vice versa) but records and/or modifies inspected traffic.

If the censor has no direct control of the DNS server but has enough visibility over DNS requests, it can also attempt a **DNS injection attacks on the side** [2, 99]. In this case, when the censor notices a DNS request sent from a client's resolver to a DNS server, it might reply faster than the uncensored DNS server, especially if that server needs to forward the query to DNS root or TLD servers [2]. The client would accept the quickest response, resulting in the adoption of the fake record. Such an attack has been reported as being part of China's UDP-based DNS

filtering, where forged responses for censored domains are injected before the legitimate ones arrive [89]. Finally, in case a DNS transfer is conducted through TCP (e.g., because it exceeds the maximum allowed size of UDP-based DNS packets), an on-path censor could also **inject TCP RST segments** [99].

5.5. HTTP-based Methodology and Web-content Filtering

HTTP-based filtering was reported to be the most detected censorship mechanism in an analysis by Master and Garman published in 2023 [124].

When the censor controls the target, it can process **header and payload filtering**, i.e., block packets of connections that contain undesired keywords/content, such as **Host**: parameters [82]. In case of unencrypted HTTP connections, the censor has multiple options: if the censor *has control over the server*, it can **influence the transfer of content hosted at the server** (e.g., modify the HTML content during transmission [103]), it can also **alter the behavior of the server**, and it can **record the exchanged traffic** between client and server for further analysis. In case of behavioral alternations, the censor can return a 30x HTTP response code that represents a **redirect** [2], i.e., the client is redirected to some error page, blocking page or a fake copy of a website where the client is brought in direct contact with the analysis backend. The censor can also **pretend that a website does not exist or that access to the website is forbidden** (response codes 404 and 403, respectively) [2, 103]. Further, the censor could **pretend that there is an internal server error** (response code 500).

Instead, if the censor has *no* control of the target server but is located on a gateway between the client and the server, it can **eavesdrop/monitor** the connection and can influence the content of HTTP connections by **injecting fake response codes** as described before. This also works if the end user uses *circumvention sites* (i.e., non-blocked websites that display bridge targets as proxies [18]). However, the injection needs to be conducted in a just-in-time fashion so that the faked server response arrives before the server's original response. To this end, the censor's gateway could handle the whole handshake with the client and pretend to be the server while the original packets are not forwarded to the server, at all. An alternative scenario mentioned by Aceto and Pescapé [2] is one in which the censor installs a **transparent proxy** on the path between the client and the server. This provides the advantage that a server's responses never directly reach the client.

Finally, a censor could redirect a user through a **Man-on-the-Side (MotS) attack** that was revealed to be used by the NSA [208]. The attack itself has been known since 1985 and is carried out on the HTTP and TCP levels jointly. The goal is to redirect a user who aims to visit a target website to a censor-controlled website where the user's browser is tricked to download a drive-by-exploit that infects the user with malware. The core idea is that the censor replies quicker to the user's HTTP packets than




Figure 4: Man-on-the-Side (MotS) attack. The client (i) initiates an HTTP request to the legitimate server. The censor intercepts the traffic and (ii) injects a TCP segment containing a redirect to a malicious server. The client's original request may reach the server, and the server will respond. However, the censor's response arrives earlier and is thus considered a duplicate by the client and discarded. The client then connects (iii) to the malicious server, which (iv) provides censored content and/or delivers malware.

the actual target, which requires the censor to inject a TCP segment. The original reply of the target arrives later than the injected one, and is thus considered a duplicate and discarded by the client. The injected segment contains the source address of the server and is destined to the IP of the client; it contains a redirect to the website of the attacking censor. The hosted replica is optically indistinguishable from the original target. Fig. 4 illustrates this process.

In addition to HTTP, HTTPS censorship has been discovered more often in recent years, probably rooted in the fact that the majority of HTTP traffic is now TLS encrypted [124]. When HTTPS is used, a censor can still perform a set of actions: (i) it can render HTTPS connections to undesired targets unattractive by **dropping randomly selected packets or disrupting connection establishment** through TCP RST packets [189]; (ii) it can **try to gain access to encrypted communication** [189] as discussed in Sect. 5.7. The above-mentioned **MotS attack** could also be conducted in case of HTTPS if the censor injects its TCP packet early during the connection establishment phase (the initial TCP SYN packet is still unencrypted, and the censor could spoof the whole TCP handshake and then the first packet that contains the redirect).

Censorship can also be conducted in a user-tailored manner, where users are first tracked and then handled individually. An important scenario refers to the adoption of the *HTTP Strict Transport Security* (HSTS) protocol. So-called **HSTS supercookies** can be used for censor-triggered tracking attempts [51, 184]. In HSTS, websites let a browser store a unique HSTS bit that the browser transmits to a website during future visits. For creating a supercookie, websites can further include (invisible) sub-components of other websites that also store such HSTS

bits, combining the unique value stored in the browser that can later on be used by the censor to identify the client while browsing.

5.6. Other Application-layer Protocols and Platforms

A censor's actions on this level are related to traffic, user accounts, and content. To allow for a better understanding, we decided to include selected non-purely technical aspects that provide a slightly broader view in this subsection.

Keyword Filtering. Like in case of HTTP and DNS, all unencrypted communication (including e-mail message bodies transferred via SMTP, IMAP or POP3, chat content or FTP filenames and payload [124]) can be **analyzed by content filters** (keywords or other forms of content and hash values [124, 82]). For instance, unencrypted communications have been censored early on [11]. Traffic matching such filters can be **blocked or redirected to monitoring networks**.

AI-driven payload analysis is increasingly being deployed to enhance simple filter heuristics, enabling sentiment analysis and image detection within communications [40]. For example, AI can now analyze chat logs for nuanced critiques or automatically detect images associated with protest movements. This allows censors to move beyond explicit keywords and suppress a broader range of dissenting communications.

Blocking and Throttling. Some popular services are also blocked with crude heuristics. For instance, Telegram was censored by **blocking large sets of IP addresses** associated with nodes responsible for providing the required computing, network, and storage resources. As a result, Telegram users applied several methods, including obfuscation as well as *IP hopping* so that IP addresses change quickly, which requires blocking IP addresses with massive collateral damage [62]. Another communication service (Twitter, now X) was reportedly subject to **throttling**, and a censor identified Twitter-usage through the Server Name Indication of TLS (see Sect. 5.7) when related domains were identified, e.g., *twimg.com*, *twitter.com* and *t.co*, [211].

Access Revocation and Content Deletion by Authorities. Another early attempt that can still be found in today's Internet censorship is deleting anonymous posts, performing or revoking access to Usenet, e-mail or FTP [11] or other online services, which are technical actions ordered by authorities. A related strategy of censors is to **delete or modify content** (videos, textual content, e-mails) of online platforms or servers if under their own management or request such removals from operating companies. With the rise of online gaming platforms such as *Steam* or *PSN*, they have become targets for censorship, too. Some of these censorship systems are in place to prevent addiction to online gaming platforms, while other censorship systems

aim to limit the exchange of political statements. In such settings, censorship is conducted based on age entries and IP address blocking [65]. If undesired messages are detected, the censor could also **pretend that a message was delivered** but actually delete it.

Replacement of Tools. Despite not being directly targeting the network layer, this censorship approach allows to perform a direct control over the resulting traffic. Specifically, a censor can interfere with both the development process or the distribution stage of a specific tool, in order to produce “censored” versions. In case of popular video and text chat services, a censor could block access to the particular original tools (e.g., their download sites) and **offer replacement tools** with censorship functionality already built in. This was done for *TOM-Skype* (Chinese clone of the discontinued Skype audio-video chat tool) and the instant messenger *SinaUC* [1, 2]. The Chinese e-mail and chat platforms *QQMail* and *WeChat* conduct censorship based on keywords and keyword combinations [109]. In addition, the LINE chat tool, the microblogging platform *Sina Weibo* and live streaming platforms *YY*, *9158*, *SinaShow* and *GuaGua* as well as Chinese online games face censorship [110, 108].

5.7. TLS-based Methodology

The majority of modern operating systems already provide a list of certificates signed by a trusted Certificate Authority (CA), rendering attacks on TLS-based communications challenging, especially unnoticed ones as browsers and other tools inform users about certificates that have not been signed by a trusted CA. For this reason, a censor would need to **deploy its own certificate**, e.g., for a websites’ replica on the client. This can be done through malware, but is not easy to accomplish. Through DNS cache poisoning, the client could then be redirected to the replica site hosted on a server with an IP address that is owned by the censor, and the faked certificate for the website must be used by the client at this point. Without DNS cache poisoning, the censor could still try to trick the user to visit a replica site using a domain containing homographs, i.e., a domain name with characters that *look* like the ones of the original website (e.g. **target-website.xyz**) but contain letters that are drawn from another alphabet, e.g., replacing a Latin *e* with a Greek ϵ : **target-website.xyz**.

In case of a man-in-the-middle (MitM) scenario, the censor could also conduct a **STRIPTLS attack** (also known as *STARTTLS stripping attack* and *STARTTLS downgrade attack*) in which the censor removes the STARTTLS command from packets so that the TLS handshake is not initiated [122]. Alternatively, the censor could apply a **downgrade attack** during the TLS handshake so that the client believes that the server only supports an outdated TLS/SSL version and/or less secure ciphers. However, the options for cipher suite downgrades are limited since TLS/1.3 only allows the selection of five up-to-date ciphers: four based on AES 128/256 using GCM/CCM and

SHA256/SHA384 as well as one using CHACHA20_POLY1305_SHA256) [168].

Further, a censor can determine the virtual hostname of the server if the **server name indication** (SNI) extension is used by the client [124, 39, 98, 211, 99, 82]. While the encrypted virtual hostname is sent to the HTTPS server using a **Host** string during the HTTP request, the SNI is transmitted earlier during the TLS *ClientHello* message and tells the server the desired virtual hostname in plaintext. The purpose of SNI is that the server can select the correct certificate of that virtual host [58] which is then used for encrypting the following HTTP request.

There is also an **Encrypted SNI** (ESNI). ESNI is available since TLS/1.3. As the established TLS channel cannot be used to encrypt the SNI (this would require to first exchange keys but they will be exchanged with succeeding packets, resulting in a chicken or the egg dilemma), one needs to employ a *separate* channel for exchanging a secret used to encrypt the SNI. ESNI works as follows [169, 39]: The administrator of the web server must publish a public key for the particular domain on a DNS server (typically in the form of a TXT record). The client fetches that key before connecting to the web server. The client uses the ESNI (instead of SNI) extension in the *ClientHello* message to encrypt the SNI with a symmetric key derived from the server’s public key and a key selected by the client through a hybrid public key encryption (HPKE) as defined in RFC 9180 [17]. However, a censor can still block connections that contain the ESNI extension in the *ClientHello* message (although the censor would act blindly, i.e., not knowing the particular domain that was requested). It could also **reduce the QoS** of the connection. Moreover, the censor could also first **spoof the target’s DNS key entry** and act as a MitM attacker between client and server. Further security considerations on ESNI can be found in [169].

Finally, for TLS versions before 1.3, censors can determine the target through the certificate transferred as part of the server’s response to a *ClientHello* message [82].

5.8. Multi-stage Censorship with Active Probing

As mentioned in Sect. 4.2, censorship is sometimes realized in a multi-stage fashion, in which the first stage is a performance-optimized stage that allows for scanning a massive amount of traffic in real-time. Flows demanding further inspection are then redirected into a second (or further) stages. Fig. 5 visualizes the concept. For instance, Aceto and Pescapé, reported that HTTP can be censored/monitored in a second stage dedicated to filtering [2]. In such a setup, the first stage redirects the client to a fake HTTP proxy using either DNS hijacking (client requests a blacklisted IP of a webserver) or BGP hijacking (client sends a packet to a blacklisted destination, port 80 or 8080).

However, by using only passive observations, a censor cannot necessarily state if a client is talking to a “benign” node or to a circumvention proxy (e.g., a Tor bridge or a




Figure 5: Multi-stage censorship: The first stage censoring device (i) conducts a high-level censorship decision (D) to either allow, block or forward the flow to a second stage censoring device (ii) for in-depth analysis.

VPN endpoint). To this aim, **active probing** refers to attempts that involve an active interaction with such nodes to infer information about the services they run [66]. Active probing also provides a link to related research disciplines: a similarity of active probing and network information hiding methods is that an adversary (censor) tries to infer (and optionally influence) the covert communication protocol (e.g., a proxy protocol or a covert channel protocol) [128], which optimally leads to gained information about involved parties and the content or type of communication.

In general, censors utilize a second-stage infrastructure for such probing approaches. Fifield highlights that active probing is linked to certain benefits for a censor [66]: (i) probing can be conducted asynchronously (time-delayed) to trigger observations, and (ii) the risk of false-positive classification (and consequential blocking) must be considered a collateral damage, but the overall blocking precision is improved by active probing. During active probing, a censor sends different types of probe requests, depending on the protocol. For instance, a censor could send tailored HTTP requests to a node that is suspected to be a HTTP proxy.

Case Study. China developed one of the most sophisticated monitoring and censorship systems in the world [215]. Reportedly, the *Great Firewall* (GFW) utilized active probing mechanisms already in 2010 by initiating own connections to remote nodes to infer if they provide proxy-related functionalities [66]. In [206], Wu *et al.* characterize the GFW when used to block websites and filter contents [12]. As fully encrypted protocols are critical for censorship circumvention, the GFW aims to detect and block this type of traffic.

The GFW operates using crude heuristics to exempt traffic that is unlikely to be fully encrypted, and then blocks the remaining non-exempted traffic in real time [206, 167]. The censorship system includes both passive and active components that function independently. For passive censorship, the GFW is capable of performing purely passive detection, traffic analysis, and real-time blocking

of fully encrypted traffic. In parallel, the GFW employs active probing as part of its active censorship. Both active and passive censorship are based on efficient and largely identical heuristics, with active censorship having additional packet length rules. Specifically, the GFW only probed experimenter’s connections with 200-byte random data, but not with 2- or 50-byte random data [206].

There are five specific rules that the GFW employs to exempt non-encrypted traffic from being blocked [206]: The first exemption rule (Ex 1) measures the entropy of traffic packets by counting the set bits (*popcount*). If a bit string fall between 3.4 and 4.6 bits per byte, the traffic will be blocked because it appears too similar to randomized data, which is characteristic of encrypted traffic. Ex 2 exempts traffic from blocking if the first 6 bytes of the connection fall within the printable ASCII range. Ex 3 exempts traffic if more than half of all bytes in the first packet fall within the printable ASCII range. Ex 4 exempts traffic if there are more than 20 contiguous bytes that are printable. Ex 5 exempts two popular protocols, TLS (as long as the first three bytes of the connection match the TLS *ClientHello* message) and HTTP (at least for the HTTP methods GET, PUT, POST, and HEAD), as the GFW is able to infer these protocols from the first 3-6 bytes of the client’s packet. Lange *et al.* figured out that China does not block any unencrypted HTTP/2 traffic [113].

The GFW can block traffic on all ports. It conducts **residual censorship** [24, 39, 197], meaning that the same 3-tuple (client IP, server IP, server port) will be blocked for 120 or 180 seconds, and the timer does not get reset during this period.¹ The GFW limits the number of connections it blocks residually, resulting in shorter blocking times when several connections are blocked in parallel. However, residual censorship makes it difficult for a proxy user to successfully connect once detected, as all involved connections between client and proxy are affected [206].

The goals of the blocking strategies deployed by the GFW are to mitigate false positives and reduce operational costs. This is mainly achieved by 1) only blocking specific IP ranges, such as those of popular data centers that offer proxy connections (affected IP ranges), and 2) using a probabilistic blocking strategy that blocks only 26% of the affected IP ranges [206]. This approach results in a lower true positive rate but also reduces false positives.

5.9. Limitations of Censorship Techniques

The presented censorship techniques are also characterized by a wide range of limitations. In the following, we review and discuss them. Tab. 3 summarizes the main limitations, organized on a per-layer basis.

IP-level. Aceto and Pescapé point out that an IP-level filter must be located on the path between client and

¹Another form of residual censorship is *long-lived* residual censorship that can be extended to 30 min and beyond [21].

Table 3: Summary of censorship techniques’ limitations according to literature (●: relevant, ◐: partially relevant, ○: not relevant)

	Censor Device Location- Dependence	Blocklist In- completeness	Over- blocking	Cryptographic Hurdles	State- holding	Notes
IPv4/v6	●	●	●	-	●	-
Routing	●	○	●	RPKI	○	-
TCP/UDP/ QUIC	◐	●	●	○	●	QUIC challenging to block
DNS	◐	●	◐	DNSSec, DNSCurve, DoH, DoT	○	-
HTTP(S)	◐	●	◐	Encrypted HTTPS traffic	○	HTTP smuggling, content transformations (website screenshots)
Other Appl. Protocols	◐	●	●	TLS-based protocols	◐	Replacement tools costly to develop
TLS	◐	●	●	ECH, ESNI	○	-
Multistage Censorship	◐	○	○	○	●	Obfuscation is an issue

target, which is not necessarily the case for alternative attempts [2], such as BGP-based hijacks (cf. Sect. 5.2). They further mention the fact that IP filtering can result in overblocking, e.g., when an IP of a virtual host that also serves legitimate websites is blocked [2, 60].² This is confirmed by Master and Garman in a 2023-study that has shown a decline in IP blocking [124]. Master and Garman mention three reasons for the decline: (i) blocklist maintenance of ephemeral IP addresses, (ii) collateral damage when blocking a CDN’s IP range³, and (iii) deployment of IPv6 with the resulting growth in addresses. One additional issue refers to resource limits when a high number of fragmented packets arrive: when a censorship device tries to reassemble a high number of packets of which only fragments (IP packets with the MF flag set to ‘1’) have arrived, a stateholding problem [83] arises.

Routing. BGP-based traffic interception became more challenging in recent years, especially due to the implementation of Resource PKI (RPKI) which prevents some hijacking methods [42]. While RPKI prevents attackers from announcing a targeted IP prefix as an origin, it is still possible for an attacker to maliciously announce a short AS-path through themselves towards the actual origin AS. Further, detection of BGP hijacks has been studied extensively [15, 42]. Master and Garman found in a 2023 study [124] that among 70 analyzed countries, BGP-based censorship was the least utilized among eight selected censoring methods. They believe that this might be linked to the fact that BGP announcements impact the routing of several more BGP routers outside the censor’s AS. As highlighted by Al-Musawi *et al.* [5], even sophisticated

BGP analysis may not reveal the complete picture of censorship evasion, necessitating the integration of additional data sources such as IP geolocation and shutdown records.

UDP/TCP/QUIC. Maintaining a port blocking list is difficult as targets can change their assigned ports over time and can be coupled to ephemeral IP addresses (see Sect. 5.1). Further, due to resource limits, there is a state-holding problem (like with any other stateful filter devices, such as traffic normalizers [83]), i.e., censor devices cannot keep track of an unlimited number of TCP connections’ states. Resource limits in terms of computing power and memory also limit how the censor is capable of handling inconsistent TCP retransmissions [83] (i.e., when a filtering device observes retransmissions of segments with different payloads which can be challenging when fragments are to be re-assembled for further analysis). A minor issue is the cold-start problem, i.e., the evaluation of already existing connections after bootstrapping the normalizer [83]. While UDP- and TCP-based blocking attempts are pretty common, blocking methods for QUIC are partially in their infancy. Elmenhorst *et al.* mention in a 2021-study that QUIC is not necessarily considered by some censors due to its novelty [60]. However, the authors expect a growing trend in QUIC blocking.

DNS. As with IP blocklists, maintaining DNS blocklists faces the problem of incompleteness and outdated entries. Moreover, manipulating DNS is more challenging when DNSSec or DNSCurve are used, as responses are signed by an authoritative server and thus can be verified for correctness by a client. As a result, injecting a faked response is not directly feasible for a censor [2] as the censor usually does not own the private key of that particular DNS server. However, a censor might still block packets directed to a particular DNS server [2] (DNSCurve can encrypt DNS requests, so the censor can only drop such packets blindly).

²Cf. Sect.5.7 on censoring virtual hosts while TLS is used.

³For instance, Wu *et al.* were able to determine collateral damage of 0.6% of all connections for China’s filter rules [206].

When *DNS over HTTP* (DoH) and *DNS over TLS* (DoT) are used, the packets between the client and the resolver are secured, but not necessarily the packets between the resolver and the queued nameserver(s), allowing DNS manipulations between these systems at least for on-path scenarios [98].

HTTP/HTTPS. While censoring attacks on plain HTTP are easy to accomplish and thus attractive to the censor, several of the HTTPS-based attacks are more challenging to realize (see Sect. 5.7) as deploying faked certificates requires additional steps, and already established encrypted connections cannot be decrypted by the censor. Further, if circumvention sites (that provide proxy front-ends for circumvention [18]) transfer blocked target websites in image form (screenshots), their blockage becomes non-trivial as keyword filtering cannot be applied anymore. Recent work has also shown that censorship can be circumvented through *HTTP smuggling* [135], where an HTTP request is nested inside the payload of a surrounding HTTP request (the `content-length` refers to a larger request, but the transfer encoding splits the nested from the surrounding request, and the censor is assumed to investigate only the surrounding request).

Other Application Layer Protocols. Blocking attempts like in the case of Telegram can easily bring larger collateral damage with them. Blocking and limiting is also challenging as there is a plethora of censorship circumvention methods available and since new online platforms appear quickly, requiring adjustments of censorship methodology. TLS-secured application-layer protocols need more effort and face the TLS-related challenges (see Sect. 5.7). The development of own (replacement) tools with built-in censorship is costly for the censor and only attractive to end-users if the originals are either not accessible (blocked) or the replacement tools are of similar quality.

TLS. Master and Garman mentioned that if an Encrypted Client Hello (ECH) is used, SNI-based filtering can be expected to be eliminated [124]. However, ECH is still an IETF draft [169] and may also introduce new censorship challenges. For example, the censor can determine that ECH is being used and, as such, they may choose to simply block all connections using ECH. Beyond ECH, an attempt was made to at least encrypt the SNI itself with *Encrypted SNI* (ESNI). ESNI censorship was analyzed by Chai *et al.* in 2019, revealing indeed less blocked sites as in case of using SNI. ESNI attacks are either blind or more challenging (DNS poisoning combined with a MitM transparent proxy).

Another limitation is *domain fronting*, where CDN-based proxies are used to access the target but also several legitimate websites, so that blocking such a proxy results in *unacceptable collateral damage* [68, 66]. This is essentially realized by constructing a TLS handshake that contains an allowed SNI value, but the TLS-encrypted HTTP

request contains a `Host` value that indicates a forbidden website to which traffic is finally routed [68, 82]. Xie *et al.* consider this approach a network covert channel [209]. However, it would be imaginable that censors aim to enumerate the number of devices behind CDNs (attacks like *structure analysis* [107] already allow counting QUIC devices behind load balancers in real-world CDN settings) to quantify the collateral damage and aid their decision-making process. Finally, users might omit SNI values so that censors cannot easily determine targets, while blocking such connections could lead to overblocking [82].

Multi-stage Censorship. Several methods have been implemented to render active probing and other sophisticated detection methods more challenging to a censor. One of these methods is to apply domain fronting (Sect. 5.9). Domain fronting has been used for several circumvention systems, including Tor, Psiphon, and Lantern [66, 68].

Another approach to challenge active probing is to minimize indicators of circumvention traffic, e.g., by employing *sophisticated covert channels* (steganography methods) that employ multiple hiding patterns sequentially or in parallel, e.g., through temporal or spatial modulation [202, 126], as well as *obfuscation extensions*, such as the Tor pluggable transport `obfs4` [209]. For instance, one temporal modulation is to frequently change proxy addresses, as realized by *Snowflake* [26, 66], which can be considered a form of *host-based scattering* [126]. Another form is to employ multiple flows, protocols, and circumvention methods simultaneously, like done by *StegoTorus* [199], hiding pattern hopping [202], *Turbo Tunnel* [67] and *Raceboat* [193]. However, censors also adjusted their active probing methods to identify obfuscation modules such as `obfs2` and `obfs3` [66].

6. Network-level Censorship Measurement

The previous section described censorship techniques. We will now investigate how the presence of these censorship techniques can be inferred, that is, *detected* or *measured*. Note that we neither discuss generic issues of Internet measurement as they have been, for instance, discussed by Paxson [156, 155], nor do we cover data analysis aspects, which are covered by standard textbooks, e.g., [47]. Instead, we will focus on the measurement of censorship-specific artifacts.

6.1. Measuring IPv4/IPv6-based Censorship

Measurement of IP-based censorship in a device-agnostic manner presents a significant challenge due to the inherent diversity and dynamic nature of the censorship infrastructure. Censors employ a wide array of devices with unique configurations and filtering mechanisms. In addition, the path between a client and a target server often traverses multiple ISPs, transit providers, and Content Delivery Networks (CDNs), making it difficult to pinpoint the

exact location and nature of IP-based interference. For this reason, IP-based censorship measurements often focus on specific devices or installations and their location within a network [165]. As such, this section will discuss several key approaches to measuring IP-based censorship, including (i) reachability tests, (ii) QoS degradation measurements, and (iii) techniques to identify the location of censor devices within the network.

(i) Measuring Reachability The most fundamental approach to determine the reachability of hosts is to conduct **scans of IP ranges** from within an AS that faces censorship. For instance, a volunteer residing in such an AS could send ICMPv4/v6 echo request packets either to the hosts of a whole subnet or to a list of selected hosts on the public Internet that are suspected to be blocked. The sheer fact whether a valid ICMP echo reply was received is an indicator of the reachability of an IP. However, such trivial probes are often error-prone. First, sending a high number of ICMP requests can trigger ICMP rate limiting [166, 80], which is often used to mitigate the impact of DoS attacks by suppressing ICMP response packets. ICMP rate limiting requires an intermediate router to drop a fraction of probe packets and requires the probe device to send requests with a slow rate [80] and with a repetition of packets, thus slowing down such scans. Second, the appearance of a high number of ICMP probe packets within a short fraction of time can lead to a probe device's IP being blocked residually. Detection of ICMP rate limitation during measurements can be done by determining (randomized) probe loss by comparing slow-rate test probes with faster-rate test probes [80].

Liang *et al.* developed *Pathfinder* [115], which measures the reachability of websites but with a focus on routing paths (and is thus discussed in this section). *Pathfinder* is a system that utilizes different routing paths within a country to determine whether a target is blocked on all of these routing paths. They discovered an inconsistency in several policies deployed by countries, which is commonly referred to as the **censorship inconsistency problem** [115]. This issue is rooted in the fact that different ISPs do not necessarily implement exactly the same blocking policies. Bhaskar and Pearce report that censorship inconsistencies have been recognized for several protocols [21]. Liang *et al.* argue that previous research neglected the fact that even *different* vantage points used to emit probe packets still take overlapping routes if the packets are destined to the same target (vantage point location was early on reported to potentially skew the interpretation of measurement results [156]). Their approach is to utilize vantage points in the form of residential SOCKS proxy IP servers so that they are located around the world. In addition, they deployed a set of control servers that represent targets and reside in uncensored environments. So-called clients feed a list of destination domains to the vantage points. Next, these vantage points try to reach the target servers and conduct HTTP requests, in which the **Host** header field contains an entry of the domain list

provided by the client. This enables the censor to spot the domain name in the header and can trigger a filtering reaction of the censor device. As vantage points are distributed and control servers are distributed, too, and reachable with different IPs, *Pathfinder* can measure several paths between vantage points and target systems and infer blocking policies on different censor-operated gateways. Moreover, the authors deployed measurement in a way that it covers different hosting providers so that the influence of peering-policies on the censoring behavior can be estimated [115]. Elmenhorst *et al.* used a mix of Virtual Private Server (VPS), Virtual Private Network (VPN), and personal devices of volunteers [60]. They did find that VPN and VPS used for testing often were less censored compared to the devices of the volunteers.

(ii) Measuring QoS Degradations QoS degradation can take different forms, e.g., packet loss or bandwidth throttling. It is feasible to measure the QoS decrease by comparing current with historical transmissions. For instance, a current flow's unacknowledged TCP packets that were retransmitted through the TCP reliability measures can be compared to historic flows. A higher number of retransmissions indicates more packet drops (loss). Throttling of flows can be measured similarly. Either in a historic manner (comparing throughput to the same destination for flows of the past) or for long-lasting flows so that the change of throughput of a flow can be monitored. However, an increase in packet loss and a reduction in throughput can also be the result of unintended or legitimate network changes in the network, such as legitimate alterations of routing paths.

Among other factors, Anderson examines latency, packet loss, and out-of-order delivery, and provides an example strategy to identify throttling mechanisms in Iran [9].

(iii) Identifying Censor Device Location To identify the location of a censor, a common approach is to apply **traceroute-like approaches** [2]. These approaches make use of the original datagrams returned ("quoted packet" in the ICMP error response messages) to detect the presence and potential actions of middleboxes within a network path. This quoted packet is sent due to many routers following the RFCs 792 and 1812. A discrepancy between the original sent packet and the quoted packet can indicate the presence of a censor. This works because when a probe packet's Time-To-Live (TTL) expires, a router along the path sends back an ICMP Time Exceeded message. The quoted packet feature includes a portion of the original probe packet within this ICMP message. If a censorship device modifies the packet before its TTL expires, those modifications will be reflected in the quoted packet received by the measurement tool. For example, if a censor injects a reset (RST) flag into a TCP packet to disrupt a connection, the quoted packet in the subsequent ICMP Time Exceeded message will contain the injected RST packet. By comparing the sent packet with the quoted packet at each hop, researchers can pinpoint the location where the modification occurred, thus identi-




Figure 6: Example case for `traceroute`-like censor location detection like [53]. Visualization inspired by Raman *et al.* [165]

fying the potential location of the censorship device. Detail *et al.* implement this technique with their tool named `tracebox` [53]. However, it is crucial to note that while TCP RST injection is a censorship technique, it would not typically trigger an ICMP Destination Unreachable message (type 3, code 3) from the censor itself. ICMP Destination Unreachable is usually sent in response to issues like a closed port or a non-existent host. Instead, the RST packet sent by the censor directly terminates the connection on the client side. The quoted packet revealing the RST would be contained in the ICMP Time Exceeded message from the router downstream from the censor, where the probe packet’s TTL finally expires.

Raman *et al.* builds on this concept with a more sophisticated detection technique [165]. CenTrace actively probes for censorship by **sending HTTP(S) requests with varying TTL values and content**. It aims to pinpoint the location of blocking by identifying the network hop where requests for censored domains fail, while requests for uncensored domains succeed. This active probing allows CenTrace to specifically target HTTP(S) censorship, whereas `tracebox` is more general-purpose and might detect various types of middlebox interference, not just censorship. The authors do however state, that their approach would be easily extended to other protocols such as DNS or SSH.

A simple scenario of this method is shown in Fig. 6, which is split into two experiments using two domains (a) and (b). The part of the control domain (a) represents the unchanged *baseline* where no censoring occurs. The part of the test domain (b) shows that a censor intercepts requests to the server with an RST response *in the path*. This behavior indicates that R3 acts as a censor. For the test domain, the TTL to be tested can be incremented stepwise until a RST is received.

Recently, censors started to actively detect and block IP-based location probes. The possibilities for them to do so are manifold, ranging from detecting abnormal traffic patterns associated with probing to actively disrupting the tools and techniques used for fingerprinting. Censors can

control request rates, examine network traffic for abnormal patterns, or discard certain packets to conceal their actions. They can alter ICMP responses and quoted packets to deceive `traceroute`-like tools, complicating the task of tracing network routes back to the censorship origins. Furthermore, censors use deep packet inspection, block specific signatures, and apply protocol-specific filters to detect and eliminate known fingerprinting tools. After discussing the techniques mentioned in this paragraph, Amich *et al.* present *DeResistor*, a tool that employs machine learning techniques for traffic and pattern generation that minimizes the risk of being detected [8]. They leverage an open source tool called *Geneva* [25] as an engine to find evasion strategies and in this context protects Geneva from being detected in turn.

Wampler *et al.* [195] investigated how vantage points can be suitable to evaluate whether the **version of the IP protocol** plays a major role. In more detail, bidirectional measurements performed over DNS, HTTP, and TLS traffic when served via IPv4/IPv6 exhibit substantial differences in censorship. In fact, many state-level censors fully block and filter IPv4 traffic but seem not to consider IPv6. On one hand, this could be due to the limited volume of users relying upon IPv6 (e.g., Morocco, Tanzania, Kuwait, and Turkey). On the other hand, such a behavior suggests that some countries implement blockages of IPv6 traffic within a specific ISP, thus making it harder to have a precise estimate (see, e.g., the case of Russia). Apart from providing a prime attempt at understanding the IPv6 “readiness” of the main blocking mechanisms, this investigation highlights an important aspect concerning censorship circumvention. Specifically, in states where only IPv4 is blocked, v4/v6 transitional mechanisms, e.g., 6-to-4 tunneling, may offer a way to stay undetected or unfiltered.

Apart from trying to *directly* detecting a censor, a valid strategy could be the continuous monitoring of possibly affected infrastructures. This implies the deployment of a distributed measurement infrastructure to monitor Internet censorship through active probing techniques. These probes emulate the behavior of regular users, attempting to access websites and online services while meticulously documenting their experiences. Advanced Internet censorship measurement systems, like ICLab (cf. Sect. 8), enhance this idea by simplifying both the setup and execution processes.

6.2. Measuring Routing-based Censorship

Routing-based censorship can be subtle and hard to detect, as they manipulate the fabric of Internet traffic flows.

To **detect cut-out or cut-off ASes or routes**, researchers can rely on the combination of various datasets with applied anomaly detection [16]. Using these, they can try to derive the application of censorship. Combining data from various sources, such as BGP routing tables, `traceroute` measurements, DNS responses, and even publicly available information such as IP geolocation and AS

relationships, allows researchers to corroborate findings and uncover subtle patterns that could be missed when analyzing individual datasets in isolation. For example, anomalies in BGP updates can be cross-referenced with `traceroute` data to pinpoint the location of suspicious route manipulation.

Dainotti *et al.* use BGP control plane data, data plane recordings, and a geolocation database to detect changes over time in Libya [49]. They mapped Libyan IP addresses (including those geolocated in Libya via MaxMind [125]) to ASes and prefixes, tracking prefix visibility in BGP tables, and correlating this with changes in darknet traffic from Libya. The combined analysis revealed a sequence of BGP withdrawals coinciding with traffic drops, confirming BGP-based blocking. However, other outages, outside of the BGP-blocking induced ones, suggested the additional use of packet filtering. Although `traceroute` data provided supporting evidence, its limitations hindered the fine-grained analysis of shorter outages. This multifaceted approach enabled them to expose details about the blackout's timing, methods (BGP disruption and packet filtering), and attribution to actions within Libya.

Al-Musawi *et al.* further present a comprehensive survey on BGP anomaly detection approaches using these types of datasets [4]. They categorized the techniques by approach, employed BGP features and effectiveness in detecting anomalies and pinpointing their origins. An examination of 21 major techniques led to their classification into five types: time series analysis, machine learning, statistical pattern recognition, validation of historical BGP data, and reachability checks. Their findings revealed that no single technique effectively combined real-time detection, differentiation between anomaly types, and source identification. Although some techniques excelled in one or two of these areas, none achieved all three. This highlighted the need for next-generation detection systems to incorporate a combination of approaches and leverage multiple data sources (updates to BGP, Routing Information Base (RIB) tables, Internet Routing Registry (IRR) databases, etc.) and features (AS-PATH length, prefix origin changes, volume anomalies, etc.) for more comprehensive and accurate BGP anomaly detection. Their classification framework and analysis provided a valuable contribution to understanding the strengths and weaknesses of current techniques and identifying key requirements for future research.

Historical and time-focused data can also be employed to detect BGP hijacking. Shapira *et al.* uses such datasets to develop a novel BGP hijacking detection methodology. Their approach utilizes an unsupervised deep learning model trained on large-scale datasets of BGP routing information. Specifically, they employ AS-level paths and prefix origin data derived from BGP update messages collected over extended periods [179]. Another approach was taken by Moriano *et al.* with their algorithm that detects the burstiness of BGP updates. This enabled them to detect large-scale updates of malicious intent [134]. Their

method leverages the observation that BGP announcements associated with disruptive events, like prefix hijacking, tend to occur in concentrated bursts of activity, followed by periods of relative quiet. Instead of simply looking at the volume of announcements, which can be misleading due to benign events like session resets, their algorithm analyzes the timing of these announcements. This burstiness metric, combined with the analysis of the affected prefixes, allows them to detect hijacking events with higher precision than volume-based methods.

Schlamp *et al.* [176] proposed HEAP (Hijacking Event Analysis Program), a framework for **assessing the validity of suspected BGP hijacking incidents**. Introducing a formal model for Internet routing, they classify various attack types and highlight the often-overlooked threat of subprefix hijacking (announcing more specific prefixes). HEAP utilizes Internet Routing Registries, topological analysis, and SSL/TLS certificate validation to distinguish legitimate routing events from malicious hijack attempts, thus reducing false positives in existing detection systems. In addition, the detection of BGP hijacks has been studied extensively [15, 42, 174].

While not having a direct impact on the measurement of an actual route or BGP censorship, it is worth noting that routing can be inconsistent. Bhaskar and Pearce highlight the crucial role of Equal-Cost Multi-Path routing in influencing censorship measurements. Their work demonstrates that varying source IP and port, which influence a packet's flow identifier (Flow-ID), can significantly alter the paths taken through a network and, consequently, the observed censorship results. To address this, they developed *Monocle*, a route-stable measurement and `traceroute` platform [21]. By controlling the Flow-ID parameters and therefore the route taken by packets, *Monocle* enables consistent measurement of DNS, HTTP, and HTTPS censorship, including detection of DNS manipulation, RST injection, packet drops, and block pages. This approach allows for more accurate assessment of censorship by isolating routing-induced variations from actual blocking behavior.

6.3. Measuring UDP-/TCP-/QUIC-based Censorship

The transport layer of the OSI model features censorship methods based on UDP, TCP and QUIC. These protocols can be used for censorship based on triggers on the same layer or on other/higher layers (due to interference). Therefore, these protocols can also be used to *measure* censorship on their own or other layers of the TCP/IP model.

Several approaches focus on **TCP packets with RST or FIN flags being set**. A straightforward approach is trying to establish a TCP connection to a target using a tool like `netcat` and observing the resulting traffic (whether the TCP handshake is completed, a TCP RST is received [140, 21] or a TCP FIN is received at some point in time [140]). A typical indicator for censorship also comes with the observation of duplicated TCP response packets, i.e., a censor making sure that the injected TCP RST

packet arrives before the original TCP SYN-ACK sent by the target [140]. However, Niaki *et al.* point out that it is feasible to determine if a target is down instead of censored if control nodes are also receiving TCP RST or ICMP destination unreachable responses at the same time-frame [140]. Further, Bhaskar and Pearce motivate an approach where probing and comparing control and vantage point traffic over multiple routes is essential for minimizing the chance for false conclusions on censorship [21]. For instance, RST packets can occur or be omitted due to normal network errors/packet loss. Bhaskar and Pearce suggest repeating measurements multiple times while enforcing larger (30 min) time gaps between successive probes.

Clayton *et al.* used this approach in [46] to test the GFW. They sent forged TCP-nested HTTP GET requests (through `netcat`) to web servers behind the GFW with and without “bad words” and observed the resulting TCP-RST packets.

Ensafi *et al.* propose a slightly more involved method in [61]. They show that it is possible to detect the injection of TCP-RST packets by observing the increase in IPv4 Identifier (IP-ID) values after sending TCP-SYN-ACK packets with spoofed source IPs. This method can detect if no blocking, client to server or server to client blocking is applied by analyzing the amount of which the IP-ID increased. Fig. 7 depicts a visual representation of the three possibilities. In our example, an increase of 2 indicates no blocking, an increase of 1 indicates server to client blocking and an increase of 4 indicates client to server blocking. It is to note that server-to-client blocking can overshadow client to server blocking. Thus, in case of bidirectional blocking, this measurement method will only indicate server to client blocking.

In 2023 Nourin *et al.* [145] used multiple approaches to measure the censorship methods of Turkmenistan with their *TMC* system. A simple approach similar to Clayton *et al.* works by sending out HTTP GET requests containing blocked words and examining the resulting TCP-RST packets. They also used an approach that omits the actual TCP three-way handshake and directly sends out two consecutive PSH+ACK packets. The first PSH+ACK packet containing a forbidden Host header, the second packet follows with 5 to 29 seconds of delay. The second packet will then be answered by a TCP RST packet. This behavior stems from the residual blocking functionality of the Turkmen censor.

Another indicator of TCP connection tampering are **TCP sequence number collisions** after a handshake has been completed successfully while RST and FIN flags are absent [140]. This is because if a censor’s blocking page response for HTTP is sent to the client, the target’s original response would not match the blocking page, resulting in larger/smaller TCP response packets and thus sequence numbers.

Elmenhorst *et al.* [60] initiated a connection with servers and recorded when a **timeout, reset or routing error** would occur. They categorized their artifacts into

TCP handshake timeouts, TLS handshake timeouts, QUIC handshake timeouts, connection resets and routing errors. Depending on when a connection processes an error or when a timeout occurred and whether it occurred for TLS and/or QUIC, they could gain insights on the type of censorship method. If both TLS and QUIC were blocked, it would point to IP-based censorship. If only TLS connections were blocked, the timing of the error can point towards TLS-SNI based or TCP-port-based blocking.

Even if the censorship does not occur based on TCP or UDP, these protocols can still be used to measure the censor. Polverini and Pottenger noted in [160] that the DNS responses of requests to DNS servers outside of China had **corrupted UDP checksums**. This corruption was a result of tampering with the DNS packets and can therefore be used to measure DNS based censorship.

Finally, multiple forms of **port scanning** using traditional techniques can be used for detecting blocked ports (e.g., by scanning a list of destinations for ports of interest) [2].

6.4. Measuring DNS-based Censorship

A straightforward approach for measuring DNS-based censorship is to deploy measurement units in particular regions of the world and let these try to **resolve certain DNS hostnames on a regular basis**. The responses and the changes in responses are analyzed over time, e.g., detecting if `NXDOMAIN` is returned or whether the resolved IP address refers to a blocking page or known censorship page. Some authors, e.g., Nabi [137], go one step further and **compare regional DNS responses with results from international DNS resolvers** from less censored (optimally: uncensored) networks. Nabi also found that the results from international resolvers were partially tampered [137]. Similarly, Kühner *et al.* [112] compared the DNS responses gained from open DNS resolvers with those of a ground truth gained from trusted resolvers (see also [157]). In [157], Pearce *et al.* crafted DNS queries that they obtained from a publicly available list from *Citizen Lab*, cf. [44], enriched with domain names from the Alexa Top 10,000 list. These queries were sent to geographically distributed DNS resolvers. The authors consider a DNS response as manipulated when it fulfills two criteria: (i) it must be inconsistent with control data, and (ii) the returned information must be clearly flawed (e.g., non-matching TLS certificates). This approach minimizes false positives. Another study by Berger and Shavitt [19] measures specific censorship of generative AI platforms by performing DNS queries, and found that especially Russia and China block some of these platforms.

A different approach is to **send DNS queries to any IPs within the censored network**. These IP addresses do not need to be those of DNS servers, they do not even need to be active at the time of probing. The goal here is to see if a false reply is received, one to a non-routable address, such as 127.0.0.1 or any other reserved IP address.




Figure 7: TCP censorship measurement based on Ensafi *et al.* [61]. Red indicates blocked packets.

Nourin *et al.* used such an approach in [145] to measure the DNS part of Turkmenistan’s censorship system, which injects a blocking DNS response to all forbidden queries, no matter their destination IP.

Hoang *et al.* used a similar idea in [86]. They sent DNS queries from the outside into China’s censored network using two target nodes controlled by the researchers. These hosts had *no* DNS server software installed, so any DNS replies must come from the GFW. These domains are then tested again from inside China by using other researcher controlled machines.

Similarly, Bhaskar and Pearce sent requests for DNS A resource records to IP destinations that also do *not* run DNS resolvers [20]. They inferred the presence of a censorship system if there was no response for a non-filtered target but one for a “sensitive” target as this must have been injected by the censor device, although there was no actual DNS resolver that would have sent a response.

There are also measurements of DNS tampering that aim to minimize false-positives by enhancing the observation period. Niaki *et al.* [140] monitored NXDOMAIN as well as inhouse IP address responses and compared responses received by vantage points with those received by control nodes. To avoid false-positives, only those responses were labeled as “tampering” if they have been observed for a duration of seven days while only the control nodes received globally routable addresses.

Another method for inferring DNS censorship was also introduced by Niaki *et al.* [140]: they detect if a vantage point receives **two responses for the same DNS request that belong to different ASes**. The assumption of Niaki *et al.* is that an on-path censor injects the DNS response on the fly while residing in another AS than the DNS server that was actually indicated by the destination IP. A special case reported by the same authors is when the **vantage point and the control nodes both receive addresses from different AS** while both of these addresses are globally routable. The reason for such an observation can be caused by CDN providers that direct traffic to data centers that are closer to the client [140]. However,

Niaki *et al.* observed that censors utilize a limited number of IP addresses for redirecting clients to block pages. To distinguish between false positives (such as CDNs) and censorship, they record a set of websites that resolve to a single IP address when requested from a vantage point. If more than θ of those websites resolve to IPs from different ASes when accessed from a control node, Niaki *et al.* consider these websites as being tampered. θ was optimized empirically to achieve a low false-positive rate, with $\theta = 11$ providing satisfying results [140].

In [20], Bhaskar and Pearce focus on a different aspect of DNS censorship measurement by investigating the **impact of source port and source IP address on the censorship behavior**. They discovered that varying source parameters influence the path taken by probe queries. As a next step, they evaluated the censorship behavior based on these changing routes. Their results indicate that the changes in source parameters also influence the censorship behavior. Bhaskar and Pearce found that most changes are “all or nothing”, meaning some combinations of parameters will experience full censorship while others will experience none.

Fan *et al.* introduced Wallbleed [64], a buffer over-read vulnerability in some GFWs middle boxes performing DNS injections. This bug allowed the researchers to extract up to 125 bytes from the middle boxes’ memory per request, allowing a more in-depth understanding of the GFWs inner workings. The exploit was triggered by sending malformed DNS requests with mismatched length fields for domain labels. During their two-year study, they found data that resembled x86_64 pointers and x86_64 instructions. Although the authors do not believe that these intrusions were code of the GFW. They also found out that the GFW would leak previously analyzed network traffic, proving an additional risk to user privacy. The bug that was exploited in Wallbleed has been fixed in 2024 and can no longer be used, but it shows a general idea how censorship can also be measured.

CDNs, load balancing and other methods can influence DNS resolution. For this reason, Tsai *et al.* came up with

an approach for **measuring DNS manipulation based on TLS certificates** [188]. The general idea here is to check the TLS certificates presented by the target servers because a valid TLS certificate is considered a strong indicator for a legitimate connection to the actual server. In their measurements, the authors did not encounter any valid certificates for blocked pages. Untrusted certificates with matching hostnames were a sign of blocked pages, which the authors confirmed with known block page fingerprints. Trusted certificates with mismatched hostnames were considered a sign of DNS manipulation. Untrusted certificates with mismatched hostnames were considered a strong sign of DNS manipulation. Lastly, the authors also considered invalid “control” certificates. I.e., certificates that are also invalid when resolving from an uncensored device. These certificates mostly stem from misconfigured web servers and therefore cannot be considered a sign of DNS censorship.

A recent method to enhance the detection of DNS tampering was introduced by Calle *et al.* [35]. The authors utilize OONI [71] probe data (cf. Sect. 8) in conjunction with supervised and unsupervised machine learning classification. Based on features of the OONI dataset, the models successfully learned expert-defined heuristics for known tampering patterns and even identified potential new censorship signatures. For unsupervised learning, they adopted Isolation Random Forest and the One-Class Support Vector Machine which complement each other. The authors argue that these models are beneficial for unlabeled data, as DNS manipulation ground truth data is hard to obtain. They further employ XGBoost as a supervised learning model that uses decision trees for accurate predictions.

6.5. Measuring HTTP(S) and Web-content Filtering

The most common approach to measuring HTTP-based censorship involves using tools that **mimic a standard web browsing behavior**, especially sending HTTP GET requests to target websites, and analyzing the responses.

Censorship measurement can also implement the so-called **HTTP fuzzing**. This is a technique that alternates/randomizes different types of inputs to a system. For HTTP in particular, this can yield insight into the behavior of the censor. Listing 1 sketches a toy example of the approach.

Listing 1: Fuzzing HTTP request example (inspired by [97].)

```
GET / HTTP/1.1\r\nHost: www.target.com\r\n\r\n
GEt / HTTP/1.1\r\nHost: www.target.com\r\n\r\n
get/ HTTP/1.1\r\nHost: www.target.com\r\n\r\n
```

The fuzzing requests sent to the server do not necessarily have to conform to the HTTP protocol standard, but rather are used to inspect the response behavior of a censor. HTTP fuzzers commonly utilize accessible tools such as Python scripting, `wget`, or `curl` to create the required HTTP requests. Fuzzing involves examining the entire content of the HTTP response to identify potential indicators of censorship, such as the presence of block pages,

error messages, or content substitution. An example of censorship-oriented HTTP fuzzing is provided by Jermy *et al.* with their software named *Autosonda*. Autosonda can reverse-engineer HTTP filtering rules and fingerprint censorship middleboxes by carefully crafted network traffic probing [97]. They discovered limitations of filter rules in multiple ways. Examples include introducing (unnecessarily) long target hostnames, testing multiple variants of `\r\n` character combinations and placements, and introducing keywords in the `Host` parameter.

There are additional approaches available for analyzing HTTP-served content for censorship. One method is to determine the **structure of HTML tags** inside websites, as censors tend to use the same blocking page structure on different systems and with different detailed content [140]. Niaki *et al.* created vectors that represent HTML tag structures, such as `<body><ul>...`, and could identify the typical vectors for several block pages [140]. The same authors identified **textual similarity clusters** in websites by processing the HTML structure with locality sensitive hashing (LSH) which was used to center candidate pages around clusters of block pages. This approach worked as censors tend to apply the same (or similar) phrases in their blocking pages but required manual inspection [140]. A third method by Niaki *et al.* is to analyze the **URL-to-country ratio** to identify unknown block pages. To calculate this ratio, the authors took clusters not centered around a known block page and counted the URLs that point to a page in that cluster and divided that number by the number of countries with a page in that cluster. The list of URLs was then sorted and inspected manually again. Newly discovered block pages had ratios between 1.0 and 286 [140].

Finally, a 2024-study by Lange *et al.* measured whether HTTP/2 is censored if it is unencrypted. Although unencrypted HTTP/2 is not used by any popular browser, more than 6% of the tested websites supported unencrypted HTTP/2 communication [113]. Indeed, unencrypted HTTP/2 communication can be observed and tampered with easily by a censor, but the authors found that it is not affected by blocking in both China and Iran, rendering it an option to, e.g., download circumvention software.

Beyond content analysis, researchers employ a controlled measurement technique sometimes referred to as **censorship tomography** to pinpoint censorship triggers [2]. Although this term is not formally defined in the literature and can be confused with network tomography [36], we use it here to describe the following method: This involves setting up a client within the censored network and a helper server outside, both under the control of the researcher. The client sends probes, such as HTTP requests, to the helper server, embedding potential censorship triggers, such as specific keywords, URLs, or modified headers. The helper server responds in a controlled manner, sending normal responses when assumed triggers are absent and simulated censorship responses (block pages, errors, or altered content) when triggers are present. By analyz-

ing the behavior of the network between the client, the helper server, and optionally a target website, researchers can pinpoint the triggers. This analysis focuses on variations in routing paths, packet loss or modification, and timing differences between probes with and without triggers. For example, consistent packet loss for probes containing a specific keyword would suggest that the keyword is a censorship trigger. This controlled approach using a helper server isolates triggers more effectively than directly probing potentially censored websites.

Measurement of **HTTPS-based censorship** is significantly more challenging than that of HTTP due to the inherent encryption provided by TLS. As Aceto and Pescaé point out in their survey, the very nature of HTTPS obscures the content of communication, making traditional keyword-based detection or content analysis ineffective [2]. However, examining the TLS handshake process itself can reveal potential triggers for censorship [145]. Censors might block specific cipher suites, target the Server Name Indication (SNI) field, or inject invalid certificates. Please refer to Sect. 6.7 for TLS-based censorship measurement.

It is also feasible to conduct **protocol-overlapping comparisons** (e.g., HTTPS over TCP vs. HTTPS over QUIC). The introduction and adoption of HTTP3 or QUIC promises many benefits regarding Internet censorship, in particular due to its integrated encryption mechanisms. However, Elmenhorst *et al.* found that, despite its novelty, this protocol is already filtered and censored in some networks [60]. They detected these censors by conducting HTTPS-based requests and subsequently redid them using QUIC. Distinguishing the response behavior made it possible to conclude whether or not a QUIC censor is involved and if one of the protocols is censored more strictly than the other.

Finally, **cache proxies** might be deployed by Internet Service Providers (ISPs) to reduce network latency for clients. Such cache proxies can infer with censorship measurements, as vantage systems might receive a reply from such a cache proxy instead of the original target. Jin *et al.* describe a method to detect such cache proxies [99]: The authors set up an authoritative DNS server for a reference domain as well as an HTTP service for that domain. The idea is that this web service provides a different website than another web service controlled by the authors (called *disguiser* server). Vantage points are then used to retrieve the website using the authoritative DNS server for resolving the IP. If a cache proxy is present, it will cache the DNS and/or web content of that connection. Afterward, the vantage point conducts a cache probe by trying to retrieve the webpage from the disguiser server (that would reply with a different response than the actual webserver of the domain). If the response is the one of the original webserver (instead of the disguiser-provided response), a cache proxy is considered to be involved. Note that this method might fail to detect cache proxies if a cache proxy is configured to only cache specific (popular) websites [99].

Notes. To omit redundant coverage, we also refer the reader to Sect. 6.1 where we explain an approach developed by Liang *et al.* called *Pathfinder* [115], which measures the reachability of HTTP sites based on routing paths. Further, we already discussed how HTTP website censoring can be identified by sending HTTP GET requests containing **Host** header fields or potentially blocked keywords [145] while monitoring the resulting RST and FIN flags in TCP response packets as well as TCP timeouts in Sect. 6.3. Similarly, we discussed the detection of residual HTTP-based censorship using TCP flags in that section. For this reason, we will not cover these strategies redundantly.

6.6. Measuring Censorship of Other Application-layer Protocols and Platforms

We can broadly distribute the remaining application layer protocols into four distinct categories: (a) communications, (b) multimedia, (c) file sharing, (d) gaming and (e) miscellaneous. In the following, we will briefly elaborate on detections of censorship on these protocols via short examples.

In the domain of (a) communications, we include instant messaging (e.g., *WhatsApp*, *Telegram*, *Signal*, VoIP as well as textual social media platforms such as *X* formerly known as *Twitter*). A typical attempt to detect censorship of these services is based on sending **IP probes** to the affected servers. For example, Ermoshina *et al.* discusses the blocking of Telegram, a widely used chat application, in Russia [62]. As the attempt to block this application was mostly based on IP addresses that used the chat protocol (and resulted in massive collateral damage), the detection of censorship was trivial [62]. A recent work by Lipphardt *et al.* [116] investigates censorship techniques deployed in Saudi Arabia and United Arab Emirates against VoIP communications. In more detail, the work analyzes conferencing tools like Zoom and Google Meet, as well as A/V services offered by IM platforms like Telegram and WhatsApp. The investigation found that detection can be done by evaluating whether Session Traversal Utilities for Network Address Translators (STUN) traffic has been selectively blocked. In fact, inspecting **dropped STUN** requests revealed the presence of middleboxes capable of filtering traversal traffic needed for point-to-point conversations, without impacting other companion features like text and media sharing.

Another approach is to **monitor the throughput** of a service over time to detect a decline that could be caused by censor-induced throttling. Such a case is presented by Xue *et al.*, which set up a measurement platform that continuously compared the available bandwidth of Twitter domains (e.g., *t.co* or *twimg.com*) with non-Twitter domains [211]. The authors operated a measurement node in the US and eight vantage points in Russia. To measure throttling, they applied a “record and reply” method that was proposed by Kakhki *et al.* [100]: first, unthrottled traffic is recorded using control nodes; second, the traffic is replied by vantage points. Each vantage point

first sends the reply traffic with sensitive TLS SNI values to the measurement node of the researchers. This is done to cause throttling. Afterwards, the vantage points send the same traffic (but with bit-flipped payload) so that throttling is *not* caused. When the vantage points send traffic to the measurement server, the server injects TLS ClientHello messages with sensitive **SNI values** and modulates the TTL values of the injected packets. This is a **traceroute-like** approach, where packets have been injected using NFQueue. By decrementing TTL values for succeeding injections, authors were able to narrow down throttling agents' locations in Russia [211]. They discovered that throttling limited the throughput to around 130-150 kbps and that seven of their eight vantage points faced throttling. They also discovered some collateral damage in the filter policies, e.g., the censor would also throttle domains such as `throttletwitter.com` as it would match `*twitter.com`.

Video streaming and multimedia (b) platforms such as *YouTube*, *NetFlix*, *Spotify* and *Twitch* as well as VC tools such as *Zoom* or *Facetime* are also subject to excessive censorship in certain regions. Note that censorship of such platforms often comes in the form of authority-driven content takedowns rather than a pure technical filtering. Aceto *et al.* highlight a procedure to detect video content that disappears over time: an MIT project called *YouTomb* **documented previously existing videos** and determine their current availability [2]. In other words, while a government order to remove content is not directly observable, disappearing online content can be identified by low-frequent polling for the content's existence. However, such attempts can result in false-positives since videos can also disappear when users behind YouTube channels decide to delete their videos, e.g., because they become outdated. Moreover, the measurement strategies mentioned above (a) can be used (sending probe packets to servers and monitoring throughput).

Filesharing (c) protocols are often used to distribute and trade illegal and legal files of all kinds. Many countries therefore have a high incentive to limit the usage of these kinds of application. However, since these services often utilize a peer-to-peer architecture, censorship measurement of the involved protocols is particularly challenging and usually involves sending probe packets. One better way of measuring its censorship is to **participate as a peer**. Dischinger *et al.* [56] describe one such approach: Their method involves a purpose-built tool named *BTTest* that synthetically generates BitTorrent flows and measures whether those flows are aborted for no apparent reason. One might also measure throughput.

Regarding (d) (online gaming platforms), almost no dedicated network-level approaches have been developed. A measurement strategy could be to **send potentially censored keywords to gaming platform's chat services while monitoring the appearance of connection resets and throttling**. In contrast to such approaches, Feng *et al.* conducted a **user study** based on

surveys in China asking participants about online gaming censorship [65]. They discovered that most participants are capable of identifying occurring censorship, but also that most are able to circumvent those in various ways, e.g., by using VPNs and by faking their ID number or the "age" entry used in their accounts [65].

Some miscellaneous protocols (e) have also been used to measure censorship. Among these are **UDP-based Echo and Discard services**, as they are used by the Censored Planet project (cf. Sect. 8) checks for keyword-based blocking using the Echo service. The Discard service is used similarly, but, as payload is only transferred into one direction, it can be used to measure whether the direction of a flow influences the blocking.

6.7. Measuring TLS-based Censorship

One example where **TLS is used to measure censorship of another protocol** was already mentioned above in Sect. 6.4 (Tsai *et al.* [188] evaluated the validity of TLS certificates as part of their DNS censorship measurement). However, in this section, we will focus on methods that use TLS to measure censorship.

While the usage of HTTPS limits censorship, triggers based on the **SNI of the TLS ClientHello message** are still analyzable by censors (see Sect. 5.7) and are thus considered for measurement. Chai *et al.* [39] measured SNI-based filtering capabilities of the GFW. To this end, they exploited the bidirectional nature of the GFW by initiating TLS handshakes from outside of China to a server inside of China under their control to trigger censorship responses. They discovered that certain SNI values would result in TCP RST packets from the GFW, indicating censorship. Additionally, they were able to verify that the GFW is stateful, as the SNI-based filtering would only trigger when a 3-way-handshake was performed before the offending TLS ClientHello message was sent. The method could be applied to measure the statefulness of other censors as well.

Chai *et al.* also investigated **Encrypted SNI (ESNI)** censorship. They used Firefox 64.0 and controlled it with the Selenium Python library to check if pages would support ESNI when accessed from the US and then when accessed from a potentially censored vantage point. In their tests, they did not encounter any ESNI based censorship.

While researching new circumvention methods for SNI-based censorship, Niere *et al.* discovered previously unseen behavior of the GFW, which they attribute to three different censorship middleboxes [141]. Two of these middleboxes have been described in literature before, the last one was not described yet. In their research the authors applied different manipulations to TLS packets like fragmentation, header field length manipulation, domain fronting or SNI removal with the goal of stopping the GFW from analyzing the SNI and therefore circumventing the censorship. They found that domains would be censored differently, depending on the circumvention method they employed. The requests were censored with either three TCP

RST packets with residual blocking, a single RST without residual blocking or a new approach with a single RST and residual blocking. By combining various TLS manipulations, the authors were able to individually target each of the middleboxes by exploiting differences in their TLS parsing behavior. This allowed a deeper understanding of the inner workings of the GFW by highlighting the existence of a third middlebox and also by targeting specific parts of the GFW for measurements.

Another method to measure censorship is to **check whether STRIPTLS attacks or downgrade attacks are conducted**, as those are often connected to censorship approaches. While not directly focusing on censorship detection, Nikiforakis *et al.* proposed *HProxy* [142], a client-side tool to detect STRIPTLS attacks. The tool works by creating profiles of the TLS behavior of previously visited sites and comparing them to the current request and responses. This works when such an attack is newly engaged, as the tool must learn the correct (attack-free) behavior first. In general, STRIPTLS as well as downgrade attacks, can be measured by comparing handshakes with targets conducted by control nodes against those conducted by vantage points.

In [92], Holz *et al.* proposed *Crossbear*, a tool for SSL/TLS MitM detection and localization. The general idea here is to have a multitude of so called “hunters” in the form of Firefox add-ons or standalone tools that **connect to web servers from multiple vantage points and collect TLS certificates**. These certificates are then compared to a stored version of the certificate. If a discrepancy is detected, the position of the attacker can be approximated by intersecting recorded routes of different hunters with and without detected tampering. Filasto and Appelbaum mention using *Crossbear* in [71] to detect content blocking.

As mentioned in Sect. 6.3 and Sect. 6.5, Elmenhorst *et al.* [60] compared HTTP/TLS based censorship with QUIC-based censorship and were able to infer if a communication was blocked on the IP, transport or TLS-level. For the TLS part, they focused on the most common ways a connection could be interrupted: a timeout during the TLS handshake or a connection reset during the handshake.

6.8. Measuring Censor-side Active Probing Infrastructure

Fifield reported that measurements of active probing infrastructure have been conducted by analyzing source IP addresses of active probing systems and analyzing the received probing packets and their overlaps in terms of payload [66]. It turned out that, despite a large number of probing IPs were discovered, these were *managed by only a few underlying processes*, which was concluded due to shared patterns in metadata (TCP ISNs and timestamps). Connecting to the identified source IP addresses led to no reaction (except a few systems which were responsive but appeared benign) [66].

6.9. Limitations of Measurement Methodology

As presented in Section 6.1 to Section 6.8, each layer of the TCP/IP protocol architecture can be targeted by a censor. However, not all censorship approaches exhibit the same effectiveness, especially due to the different complexity of the involved protocols and use cases. This also reflects in the correctness and quality of the related measurements methods. To provide an overall view, Tab. 4 summarizes the main cons of observed for each major protocols. As it can be seen, probing device location and genuine network malfunctions are an issue common for all the considered protocol/layers. For the sake of brevity, in the following, we solely address the most relevant limitations.

Measurement Level	Vantage Point Location-dependence	Network Malfunctions/Routing	CDNs, load balancing, proxies, NAT	Dataset quality
IP	●	●	●	●
Routing	●	●	●	●
TCP/UDP, QUIC	●	●	○	○
DNS	●	●	●	●
HTTP(S)	●	●	●	●
Other Appl.	●	●	●	●
TLS	●	●	○	○
Multi-stage Cens.	●	●	●	●

Table 4: Summary of measurement methods’ limitations according to literature (●: relevant, ●: partially relevant, ○: not relevant).

IP/IPv6. Although IP-based censorship measurement provides valuable information, its effectiveness is inherently limited by the dynamic and often opaque nature of today’s Internet architecture. The path traversed by sent packets cannot be ensured by the sender, leaving room for changes in routing paths for packets even of the same flow. This challenges the detection of blocking attempts on particular routes [115]. The increasing reliance on CDNs makes it difficult to pinpoint the exact location of blocking, as a single IP address might serve content for multiple domains across geographically diverse servers. Similarly, the widespread use of Network Address Translation (NAT) obscures the true origin of connections, hindering efforts to attribute censorship to specific actors or networks. NAT is commonly used to allow multiple devices within a private network to share a few public IP addresses. This is particularly relevant in scenarios where the client performing the measurement is behind a NAT device, as it makes it appear as if the connection originates from one of the NAT gateway’s public IPs, but the client cannot necessarily control the use of a particular IP and its consistent use over multiple probing sessions. Load balancing techniques further complicate the measurement by distributing traffic across multiple servers, making it challenging to discern whether inconsistencies are due to censorship or legitimate

network management. Moreover, the accuracy of IP geolocation databases, often used to map IP addresses to geographic locations, can be unreliable, especially for shared or dynamically assigned IP addresses, leading to the potential misattribution of blocking events [5]. These limitations underscore the need for more sophisticated measurement strategies that combine IP-level analysis with deeper packet inspection, side-channel observations, and data-driven approaches to obtain a more complete understanding of censorship practices. As highlighted by Al-Musawi *et al.* [5] even when applied to BGP manipulation, the need for constant vigilance in cross-referencing data is important.

Routing. The distinction of malicious activity from an error can be difficult. As Cho *et al.* report that simple human error can and has been the cause of routing anomalies [42]. They give an example of an operator mistyping an 8 instead of a 9 in his own prefix. Moreover, pinpointing the actual censor can be hard, as he might be several hops away from the actually affected network. Further, as many detections make use of datasets, the detection can only be as good as the quality of these datasets. Acquiring up-to-date, sensible and extensive datasets can be challenging. As in case of detecting IP-based blocking, routing paths for probe traffic usually cannot be ensured [115], i.e., a test can barely fix a set of particular routers to be traversed in a given order.⁴ While Monocle [21] offers a significant advancement in route-stable measurement, challenges remain in fully understanding and quantifying routing-based censorship. The dynamic nature of routing, where paths can shift due to network conditions or traffic engineering, makes it difficult to capture a complete and static view of censorship infrastructure.

TCP/UDP/QUIC. Depending on what kind of measurement techniques are used, different limitations apply. Timeout failures can be an indicator of censorship, but it can also be the result of genuine network malfunctions, and thus, measurements are repeated regularly to observe timeouts over a longer time [60]. Moreover, duplicated TCP response packets can also be caused by legitimate reasons, such as due to HTTP load balancers [140], and even if vantage points and control nodes do not gain matching responses after sending initial SYN packets to a target, this could still be caused by network errors and can be considered as “uncertain” [140]. Methods like proposed by Ensafi *et al.* [61] that require specific setups of client and server behavior (global IP-ID on the client, an open port on the server) obviously depend on such circumstances. Another, more general, limitation is the directionality of the censorship method. For instance, Nourin *et al.* [145] report that the Turkmen firewall is bidirectional.

⁴Even if this could be changed by using IP source routing, a censor could still ignore the source route setting. Moreover, such packets might be considered suspicious and would potentially be dropped by intermediate routers.

DNS. As pointed out in [48, 140], detecting DNS manipulation faces issues due to CDNs, content localization and load balancing as control measurements might become outdated quickly or are routed to regional data centers. As the same domain might be resolved to different IP addresses based on time of day or probe location, it might be difficult to decide if a mismatch in probe results stems from censorship or legitimate reasons. This is why we recommend the additional evaluation of TLS certificates as done by Tsai *et al.* [188] and described above.

A different limitation that was described by Nourin *et al.* in [145] is that a censor might learn that a certain host is used as a measurement probe and suspend the censoring for this host. The authors noticed this with their DNS measurements, but the same idea can be applied to other measurement techniques as well. This is similar to malware detecting the presence of a debugger or virtual machine and suspending any malicious behavior.

Further, Tang *et al.* [185] show that the generation of domain probe lists, which was often crowd-sourced and manual, can be improved by automation. High-quality domain probe lists are not only essential to gain useful measurement results but also require continuous adjustments.

HTTP(S). Websites frequently use dynamic content, making it difficult to establish a stable baseline for comparison. Variations in timestamps, personalized elements, and server-side updates can trigger false positives, mistaking legitimate changes for censorship. For HTTPS, the encryption poses the biggest challenge for measurement. The inspection of content is infeasible without sophisticated techniques, requiring a reliance on metadata, handshake analysis, and side-channel observations (also cf. Sect. 6.7).

Other Application Layer Protocols and Platforms. Measuring online platforms’ and application layer protocols’ censorship is often driven by user-agents that implement the particular protocols and utilize APIs of the platforms. Closed platforms cannot be inspected with automated tools or – alternatively – would require massive reverse engineering effort to develop measurement tools. For this reason, measuring every single service must be considered a costly effort, especially since online platforms’ popularity changes quickly and new platforms emerge regularly. Further, geo-blocking must be considered, e.g., through geo-diverse vantage points. Finally, the participation as a *peer*, like in case of *BTTest*, must mimic typical user behavior to gain high-quality measurement results and often requires manual account registration.

TLS. A general issue, that is mentioned in works discussing TLS-based measurements, is the small amount of vantage points used for a measurement [39, 188]. Additional geolocation-diverse vantage points could lead to a more complete view of the censorship methods, as some censors differentiate based on the source or destination of

packets and in some cases other routes to the target can result in different censorship behavior.

Measurements based on the validity of TLS certificates could be falsified by state-level actors employing rogue CAs that find their way into the trust stores of mainstream browsers like Chrome or Firefox. Researchers have to rely on the browser developers to keep those trust stores clean [188], as such rogue CAs could falsify the reference data.

Given the above approaches, each provides an own vector of insight (e.g., faked TLS certificate or presence of STRIPTLS attacks). It would be beneficial to *combine* the ideas of existing TLS-based measurements to overcome this limitation. For instance, a large set of geo-diverse vantage points and control nodes could be used to compare as many SNI values as possible while at the same time TLS handshake results are recorded and compared to detect downgrade attacks and further try to detect STRIPTLS attacks.

7. Censorship and Censorship Measurement of Circumvention Tools

In this section, we will first introduce censorship techniques for different types of circumvention tools. Afterward, we explain the methodology that can be used to measure censorship of circumvention tools. Following that, we discuss limitations of these censorship and measurement techniques.

7.1. Censorship of Circumvention Tools

In this section, we discuss the censorship of VPNs and other circumvention tools, such as Tor, I2P and Psiphone. Since the techniques for different tools do overlap, we decided to cover them together. We concentrate on the most important censorship methods and refer the reader to specific surveys and papers when it comes to related topics, e.g., (Tor-)deanonymization attacks [101], surveillance aspects, analyses of circumvention-related court case [187], or surveys on censorship *circumvention* itself [103, 190].

Covered circumvention approaches: In addition to classic VPNs, there is a set of additional circumvention tools that are not pure VPN solutions. Such tools either (i) provide a cryptographic transmission that (like VPNs) provides an encrypted tunnel, optionally with sender/receiver anonymity, (ii) obfuscate the traffic, (iii) hide the traffic steganographically through covert channels, or (iv) combine two or three of these ideas. Over multiple decades, *The Onion Router* (Tor) [55, 171] has emerged as a widely used tool to enhance online privacy and bypass censorship. Its decentralized network architecture, which routes traffic through multiple layers of encryption (hence the “onion” analogy), makes it difficult to trace user activity and block access to websites [55, 190, 171, 22]. Tor can be enhanced with modules that provide traffic obfuscation or steganographic enhancements. In addition to the main Tor imple-

mentation, other variants of Tor exist, e.g., *Arti*⁵, which is a Rust-based implementation. However, in this article, we will not distinguish between Tor implementations. We further subsume Tor-related enhancements in this section, such as proxy services for non-onion services, e.g., *Onion-spray*⁶ and *Oniongroove*⁷. Hyphnet (also called *Freenet*) as well as its fork I2P (*Invisible Internet Project*) operate on a mix-network structure similar to onion routing. Both tools provide several enhancements and features. Another popular censorship circumvention tool is *Psiphone*⁸, which supports several different features to bypass censorship, including VPN functionality with tunnels, utilization of proxy servers, and traffic obfuscation.

Blocking Access to well-known Services. One general idea for censoring circumvention tools is to block the access to well-known entry points of those tools.

As VPNs utilize a centralized setup, it is easy to see that VPN endpoints, the servers to which VPN clients connect, are frequently targeted by censors and network operators looking to restrict VPN usage [203]. **Blocking VPN endpoints** already stops the initial creation of VPN tunnels that are used to bypass censorship or access restricted content. Common blocking methods include **IP address blacklisting and port filtering**, where the censor maintains a list of known IPs and ports of the VPN server and blocks connections to them [73]. This blocking approach can be countered by VPN providers simply shifting to new IP addresses, leading to a constant cat-and-mouse game. This aspect also affects other circumvention tools with well-known, central servers. Censors might also use active traffic analysis techniques to identify VPN entry servers. Identifying VPN endpoints can be done through **active probing** [66], where censors send connection requests to suspected VPN servers and analyze the responses (or lack thereof).

A specific detail of Tor are **directory servers**. They are crucial for bootstrapping a connection into the Tor network, and thus a frequent target of censorship. Censors can block access to these servers by blacklisting their IP addresses or domain names, preventing clients from retrieving directory information necessary to build Tor circuits. Winter [203] observed that the GFW blocks access to most Tor directory authorities by filtering their IP addresses, thereby preventing direct access to the Tor network from within China. This necessitates the use of Tor bridges (i.e., Tor relays that are not listed publicly) or alternative methods to bootstrap into the network, and highlights the importance of diverse entry points for circumventing censorship. However, a censor could also try

⁵<https://tpo.pages.torproject.net/core/arti/>

⁶<https://onionservices.torproject.org/apps/web/onionspray/>

⁷<https://onionservices.torproject.org/apps/web/oniongroove/>

⁸<https://psiphon.ca/th/index.html>

to publicly announce malicious Tor bridges through the very same channels. Additional potential attacks against directory servers have been described early on [55], including not only protocol-level attacks but also causing distrust to cause directory server dissent.

Traffic-based Blocking. Instead of blocking the access to a circumvention service based on lists of known entry points, it is also possible for censors to detect circumvention tools by analyzing and fingerprinting their traffic. This can be used to block the newfound entry points or to drop/disrupt active connections.

Such tactics can involve **Deep Packet Inspection (DPI) and VPN fingerprinting**. By analyzing packet headers and payloads, and aggregating packets into flows, DPI systems can identify VPN protocols (like OpenVPN or WireGuard) and block or throttle the corresponding traffic, even if it is encrypted [82]. For example, a DPI system might identify OpenVPN traffic by looking for specific opcodes (used to indicate the message type) or packet sizes within the encrypted TLS control channel. As Xue *et al.* demonstrate, even “obfuscated” VPN services, designed to evade detection, can be fingerprinted and blocked through these methods, highlighting the ongoing challenge of maintaining unrestricted access to VPN technology [210]. They hereby use a two-phase framework. First, the censor passively fingerprints flows by analyzing OpenVPN opcode sequences and OpenVPN ACK packet patterns within the initial handshake, exploiting OpenVPN’s predictable structure even with encrypted payloads. Suspect flows are then actively probed, triggering unique timeout behaviors in OpenVPN servers due to their packet processing and handling of invalid lengths. These timeouts act as a side channel, confirming OpenVPN’s presence even with probe resistance enabled. This combined approach achieves high accuracy with low false positives, effectively identifying encrypted VPN traffic even in the presence of common obfuscation techniques.

Gao *et al.* [73] propose a statistical method for VPN traffic classification based on Payload Length Sequence (PLS). First, they analyze the payload length distribution of various VPN protocols and web browsing traffic. They observe that different VPNs exhibit distinct payload length patterns due to variations in protocol design and obfuscation techniques. Based on this, they construct a PLS feature vector by extracting the lengths of the first n packets in a flow, excluding retransmitted packets and zero-length payloads. The value of n is determined experimentally. This PLS vector captures the sequential order of payload lengths, which is a distinctive characteristic of different VPN protocols. They then use machine learning algorithms to train classifiers on a labeled dataset of VPN and web browsing traffic. The classifiers learn to associate specific PLS patterns with different VPN protocols, enabling them to classify new, unseen traffic flows. Their results show that Random Forest achieves the highest accuracy, reaching 98.68% when using a PLS vector

of length 6, indicating that this statistical approach based on payload length sequences is effective for VPN traffic classification, even for obfuscated VPNs.

Similar to the methods mentioned above, Saputra *et al.* demonstrate the application of **DPI to detect and block Tor traffic**, mirroring techniques used for VPN censorship [175]. Their method analyzes the TLS handshake within Tor connections, focusing on characteristics like cipher suites and server name indication (SNI) to distinguish Tor traffic from regular HTTPS. While effective against older Tor versions, this approach is less successful against newer iterations employing obfuscation techniques like ScrambleSuit (cf. [204]). Similar to VPN censorship, **active probing of suspected Tor bridges** could complement DPI-based detection, enabling censors to confirm bridge relay status and subsequently block them. Their proposed blocking mechanism combines DPI with IP blocking by extracting destination IP addresses from detected Tor traffic and adding them to proxy server blocklists. Ultimately, Saputra *et al.*’s work [175] highlights the parallels between Tor and VPN censorship, showcasing how DPI and active probing are employed to target both, while also emphasizing the limitations of static detection methods in the ongoing arms race of censorship and circumvention.

Moreover, attackers can leverage traffic analysis to observe communication patterns between nodes and the rendezvous point, potentially disrupting these interactions, as stated by Karunanayake *et al.* [101]. They further state that compromising the rendezvous point itself allows attackers to intercept, modify, or block communication, effectively disrupting network operations.

Wails *et al.* [194] developed and evaluated a new approach for circumvention tool traffic identification. They developed a deep learning flow-based classifier that is combined with a host-based approach, collecting flows over time for each host. For their training, they used over 60 million real-world flows to more than 600,000 destinations. While the detection results for classical, solely flow-based detection remained low in a real-world scenario, the host-based approach had perfect recall and low false positives. The authors therefore raise concerns about host-based detections and argue for the development and research of circumvention systems that are resistant to such new detection approaches.

Website Fingerprinting for Circumvention Tool Detection. A special form of traffic-based blocking attempts are **website fingerprinting (WFP)** attacks. WFP aims to detect which websites a circumvention user accesses and was discussed for several circumvention tools, such as Tor [55, 154, 133, 41] and Psiphone [59].

While WFP can mostly be considered a deanonymization attack that is used when a known circumvention tool is present, it can also be used to identify if a circumvention tool is used in the first place. Ejeta and Kim [59] present a WFP attack on Psiphone: First, a censor (or a law-enforcement agency) visits the target website several

times using a circumvention tool and records the generated traffic with a capturing tool such as `tcpdump`. Next, metadata is extracted from these flows to gain feature values: packet length uniqueness (i.e., assigning a packet length a value between “1” (unique packet length) and “0” non-unique packet length), total transmission (total incoming and outgoing packet count and total connection time), length of the first 20 packets of a connection, packet order (total number of packets before the current packet of a sequence as well as the total number of incoming packets between an outgoing packet and its predecessor), concentration of outgoing packets (“number of outgoing packets in non-overlapping spans of 30 packets”), and bursts [59]. These metadata are used to generate a *fingerprint* model of the particular website. Afterward, the censor needs to record the (assumed) traffic of an end-user visiting the same target and extract metadata from that flow as well to form another fingerprint used for comparison. Due to both traffic recordings never being an exact match (e.g., because of different encryption keys, network congestion etc.), Ejeta and Kim [59] used kNN and SVM classifiers to determine if both fingerprints match. The reported success rate was strongly dependent on the circumvention tool used. They gained good classification results for detecting *Psiphon* traffic but reached a low accuracy when they tried to identify visited website targets among a set of previously chosen ones.

Detecting Steganography and Obfuscation Characteristics. Finally, there are many methods to detect, limit and prevent censorship circumvention based on steganography and traffic obfuscation. Such circumvention methods are applied in a way such that Tor or VPN traffic is encapsulated into a steganographic channel (serving as a tunnel) to let the traffic appear harder to detect. Similarly, traffic obfuscation as well as the replacement of previously recorded benign traffic [118] can be used to let the circumvention traffic appear as legitimate traffic of another type (e.g., circumvention tool traffic could appear as legitimate web traffic if that is not blocked). To this end, many detection methods exist, both based on statistical methods and based on machine learning methods [128, 212, 202, 37, 201]. Further, attempts to block or limit such channels (decrease QoS or limit channel capacity) are often forms of traffic normalization [128, 202, 201]. For instance, traffic normalizers might remove or overwrite unused/reserved header fields to block simple forms of covert channels in these fields or slightly delay packets to limit the capacity of timing-based channels. To this end, the available methods consider a broad set of traffic features, such as packet lengths, inter-packet times, TCP and IPSec packet sequences, packet header field values, and the structure of dynamically composed packet headers [202, 128, 37, 201]—just to mention a few.

7.2. Measuring Censorship of Circumvention Tools

In the following, we discuss different methods to measure the censorship of circumvention tools. Again, as these methods overlap, we cover them together.

A particularly interesting target to censor are different types of **directory servers**, as these are essential for bootstrapping clients into a network. Examples include Tor directory servers and I2P reseed servers.⁹ Iszaevich, for example, presents a **user study** conducted in Mexico, where he analyzed whether the Tor **directory authorities** (DirAuths) are actively censored. To do so, he **distributed a questionnaire and a guide for traceroute-based measurements** [96] and found that a significant amount of DirAuths are censored on an IP level by censors.

In addition, Dunna *et al.* performed a study in which they **placed multiple previously unpublished Tor bridges** in the US, Canada, and Europe to connect to them through Virtual Private Servers (VPS) deployed in Chinese clouds [57]. Within these VPS, they deployed Tor clients to connect to said Tor bridges. For probing connection initiation in a lightweight fashion, they additionally used the *Tor Connection Initiation Simulator* (TCIS) to quickly derive the necessary metrics for their analysis (e.g., duration of blocking). They found that China blocks both unpublished and published relays. Moreover, the gained insights on the blocking allowed them to derive circumvention strategies [57]. Such active probing is a common and effective strategy to detect censorship for both, VPN tools (cf. Xue *et al.* [210] on OpenVPN) and Tor-like tools, as it can reveal blocking through connection failures or timeouts. Network measurement platforms such as OONI [71] provide valuable data to detect Tor censorship by collecting measurements from geographically diverse vantage points.

Complementing Tor censorship measurements, Hoang *et al.* employed a novel approach using VPN Gate’s distributed network of **volunteer-run VPN servers to measure I2P blocking** globally [88]. Their methodology involved four key techniques: 1) resolving I2P domain names and reseed server addresses using local and open DNS resolvers to detect DNS poisoning, 2) establishing HTTPS connections to the I2P homepage over VPN tunnels and monitoring the TLS handshake for SNI-based blocking, 3) capturing and analyzing network traffic for injected TCP RST/FIN packets during access attempts to I2P resources, including their own set of I2P relays, and 4) comparing fetched HTML content with legitimate I2P webpages to identify block page injection. This combination of DNS, TLS, TCP, and HTTP measurements, deployed through geographically diverse VPN Gate servers, allowed them to detect I2P censorship in several countries, including China, Iran, Oman, Qatar, and Kuwait,

⁹I2P reseed servers are used by newly joint relays during the bootstrap procedure and are a type of relay whose role is to provide information about other I2P relays [88].

revealing various blocking techniques employed against the anonymity network.

Another aspect measurement studies focus on is the **delay between circumvention proxy appearance and a censor's blocking reaction**. For instance, VPN entry points, I2P reseed servers and Tor bridges appear at some time, but it would be valuable to gain insights into how long it takes for a certain censor to block such new entry points. A Tor bridge-focused measurement was designed by Fifield and Tsai [69, 66]. They implemented two scripts. The first script ran TCP connection attempts to Tor bridges every 20 min from China, Iran, and Kazakhstan and reported gained error codes. In parallel, a control script made the same connection from the US. The second script realized a whole Tor-in-obfs4 connection from Kazakhstan. The authors considered only “default bridges”, i.e., non-secret bridges known within the Tor community. Similar experiments could be conducted to reach proxies of other circumvention tools.

7.3. Limitations

imitations of Circumvention Tool Censorship. Blocking circumvention tool traffic, while a common censorship tactic, faces inherent limitations due to the distributed and adaptable nature of these technologies. Blocking VPN server IP addresses, as highlighted in Xue *et al.* [210], is a continuous cat-and-mouse game, as VPN providers can easily shift to new IPs or deploy obfuscation techniques. Similarly, blocking Tor relays or bridges requires constant monitoring and updating of blocklists, as new relays and bridges appear frequently. Though effective for identifying VPN protocols or Tor traffic patterns (as shown by Saputra *et al.* [175]), DPI is computationally expensive and susceptible to evasion through obfuscation or traffic manipulation techniques. In addition, correlation attacks using meta-data of traffic in form of volume or timing behaviors can be rendered difficult through obfuscation techniques, such as packet size modulation.

WFP faces several limitations that are practically relevant for censors. Generated website fingerprints can become outdated quickly, e.g., when new forum posts are added to a website hosting a discussion forum. Cherubin *et al.* highlight that several studies do not consider a real-world setting, e.g., relying on synthetic traffic and target websites that might not be visited by real Tor users (among other limitations) [41]. Similarly, Ejeta and Kim consider a scenario in which a user conducts no parallel activity while visiting a certain target website, and the censor is aware of the start and end time of a pageload [59]. Further, WFP is usually conducted on flows that do not employ any obfuscation features, WFP does not perform well on such flows [154]. Other issues, e.g., website caching behavior and lack of heterogeneous users, browsers, and systems, are covered in [133] and [41].

Furthermore, blocking access to VPN or Tor websites and documentation, while hindering discoverability, cannot prevent determined users from obtaining the software

through alternative channels. Moreover, the decentralized nature of both VPNs (various different providers) and Tor, with multiple entry and exit points, makes it difficult for censors to achieve *complete* blocking without resorting to drastic measures like nationwide Internet shutdowns, which often carry significant political and economic costs.

Limitations of Circumvention Tool Censorship Measurement. These kinds of measurements have several implications and limitations. Iszaevich states that since measurements were conducted by local volunteers, misconfigured experiments or old clients are a source of error [96]. In relation, as Dunna *et al.* also state, the biggest issue is the locality of the experiments [57]. For optimal measurements, people actually need to physically be in the country where censorship is applied. Depending on the laws of that country, this can be a significant risk. Moreover, censors actively adapt their techniques, making detection a moving target. The use of obfuscation methods, such as those employed in some VPNs or Tor's ScrambleSuit [203], can make it difficult to distinguish VPN or Tor traffic from regular HTTPS, hindering DPI-based detection: obfuscation makes it more challenging to detect censorship itself, as the DPI systems may fail to recognize the censored traffic as such, leading to false negatives in censorship measurements. It also complicates the identification of specific blocking techniques, as the obfuscated traffic might not exhibit the clear signatures associated with methods such as RST injection or block page redirection. As in case of previously covered methods, accurately measuring the impact of censorship, like throttling or blocking of specific servers, on circumvention tools can be challenging due to network fluctuations and the difficulty of establishing stable baselines. False positives and negatives are common, particularly when relying on automated measurement or limited vantage points.

8. Measurement Platforms & Datasets

Developing censorship circumvention methods requires an in-depth understanding of censorship. For this reason, the development of bypass tools must consider measurement datasets. These datasets can be used not only to tailor circumvention tools, but also to evaluate these tools. This section discusses the available censorship datasets and a comparison of these datasets. The above-mentioned datasets are current data or data from recent years.

CensoredPlanet. CensoredPlanet is a censorship measurement platform that captures and analyzes measurements from four different measurement techniques, namely Augur, Satellite/Iris [178, 157], Quack [191], and Hyperquack [164]. The goal is to complement existing datasets like OONI [150] and ICLab [140] by providing a larger scale, coverage, and duration of measurements. Until 2020 [183] they collected and published over 21 billion data points

from over 20 months of measurement operations. However, they continuously publish new data on their platform that is available online [38].

The platform captures baseline data from 2,000 domains using over 95,000 vantage points worldwide. This data is collected, aggregated and analyzed to provide a longitudinal view of censorship and allow distinguishing between localized and countrywide censors [38]. In practice, they were able to reduce false positives, modeling the data as time series and using the Mann-Kendall test. With this, they were able to detect 15 censorship events over their measurement period, two thirds of which were not previously reported.

CensoredPlanet not only supports long-running broad measurements, it also has “rapid focus capabilities”, allowing for timely measurements of specific networks. The authors used this to investigate changes in the blocking behavior of Cloudflare IPs from Turkmenistan.

CensoredPlanet measures censorship on six different protocols: IP, DNS, HTTP, HTTPS, Echo, and Discard. IP, DNS, HTTP and HTTPS were selected for their position as censorship targets. Echo is used to measure keyword-based censorship on the application layer by checking if the selected servers successfully reply to such requests containing potentially censored keywords. The Discard protocol is used to measure the directionality of the censor.

CenDtect. CenDtect utilizes the data of CensoredPlanet as input for their censorship detection model. Tsai *et al.* [189] implemented a system that uses unsupervised machine learning. They therefore implemented a clustering approach based on decision trees and DBScan. With this, they want to make the data easier to consume for other researchers, journalists or NGOs. The output of the system consists of decision trees, which describe the censorship rules and a corresponding domain list showing which domains are affected by the censorship rule. They also discuss the challenges of censorship datasets, like the lack of ground truth, the volatility of test lists and the sheer amount of data. While they do not provide a new dataset, they aim to augment existing data to make it more usable and are evaluating the possibility to directly integrate CenDtect into the CensoredPlanet platform. In consequence, we do not list CenDtect in Tab. 5.

Encore. Burnett *et al.* analyzed Internet censorship in [30], specifically in the form of web filtering, particularly focusing on its extent. To achieve this, they introduced *Encore*, which enables continuous and longitudinal measurement. A key advantage of the system is that users do not need to install custom software, emphasizing its ease of use.

Encore utilizes web browsers on Internet-connected devices as vantage points, allowing it to gather a large and diverse set of data. The major contribution of Burnett *et al.* is that meaningful conclusions about web filtering can




Figure 8: Encore measurement methodology (based on [30])

be drawn from the side channels that exist in cross-origin requests.

An origin is defined as the protocol, port, and DNS domain (“host”) [139]. Websites can send information to another origin using HTTP requests across origins. The cornerstone of Encore’s design is to use information leaked through these cross-origin embeddings to determine whether a client can successfully load resources from another origin. For example, an image embedded with the `<img>` tag can trigger an onload event. Using this kind of information, Encore can make determinations about the extent of Internet censorship. The scope of web filtering can range from individual URLs to entire domains [30].

Measuring web filtering with Encore involves three parties (Fig. 8): (i) the web client, which acts as the measurement vantage point; (ii) the measurement target, which hosts a web resource suspected to be filtered; and (iii) the origin web server, which serves a web page to the client, instructing it on how to collect measurements with tasks, which are small programs (self-contained HTML and JavaScript snippets) that attempt to load a web resource.

As of 2015, Encore had recorded 141,626 measurements from 88,260 distinct IPs across 170 countries, including China, India, Egypt, South Korea, Iran, Pakistan, Turkey, and Saudi Arabia. However, Encore has a significant limitation: it can only observe the accessibility of individual web resources and cannot determine the exact method of filtering. This limitation positions Encore as a supplement to other censorship measurement systems, which are capable of more in-depth analyses but encounter greater challenges in deployment.

GFWatch. GFWatch is specifically tailored to measure the behavior of the GFW by identifying censored domains and detecting forged IP addresses in fake DNS responses. The GFW employs DNS injection techniques, returning

falsified responses. To summarize the function of GFWatch: Nodes inside China run DNS tests, whose results are then compared to DNS resolving results of nodes outside China to detect DNS-based filtering. This DNS poisoning not only impacts users within China but also pollutes global DNS resolvers like Google or Cloudflare. As a large-scale longitudinal measurement platform, GFWatch provides a continuously updated view of the GFW's DNS-based blocking behavior [86].

Hoang *et al.* revealed that this manipulation of DNS records has far-reaching consequences beyond China's borders [86]. To mitigate its effects, they proposed strategies to detect poisoned responses, sanitize polluted DNS records, and assist in the development of circumvention tools to bypass the GFW's censorship. To achieve its goal, GFWatch must continuously discover and test new domains as they emerge. For that, GFWatch uses top-level domain (TLD) zone files, which are updated daily. Additionally, GFWatch provides ongoing insights into the pool of forged IPs used by the GFW, helping researchers better understand and counteract China's DNS censorship tactics [86].

GFWatch probes the GFW from outside China to identify censored domains (step 1) and verifies these findings using controlled machines within China (step 2). It employs UDP-based DNS queries, as UDP is the default carrier for DNS and requires fewer network resources compared to TCP-based measurements. The procedure follows multiple steps: (i) The primary probing machine located in a US network sends DNS queries for test domains to two hosts controlled by Hoang *et al.* in China. These Chinese hosts lack DNS resolution capabilities, so any DNS responses received by the main prober originate from the GFW. Positioning the two Chinese hosts in different ASes helps address the GFW's centralized blocking policy and detect potential regional variations. (ii) GFWatch transfers detected censored domains to its Chinese hosts, which then probe these domains from within China towards the controlled machine in the US to confirm that the censorship occurs both inside and outside China. To account for potential inconsistencies, such as the GFW occasionally failing to block access under heavy load [13], GFWatch tests each domain at least three times daily [86].

In [77], Hoang *et al.* tested a total of 534 million distinct domains. Over a nine-month period, they found that 311K individual domains were filtered. To analyze the evolution of blocking rules rather than just the number of censored domains, they identified the most general form of each censored domain (the so-called base domain) that triggers censorship. As a result, they discovered 138.7K base domains. Hoang *et al.* also found that 41K censored domains were overblocked due to textual similarity with a censored base domain, despite not being subdomains. The dominant censored categories were business, pornography, and information technology [86]. They also found that forged IPs follow distinct patterns, rather than being injected entirely at random. They identified 11 groups: 10

static groups with consistent, predictable responses (e.g., `qcc.com.tw` always returned `89.31.55.106`) and one dynamic group with variable manipulations. Since the GFW frequently changes spoofed IPs within a certain range, DNS censorship can be effectively detected by comparing the returned IPs with the pool of spoofed addresses discovered by GFWatch [86].

Using this methodology, it is possible to sanitize polluted DNS records in the cache of public DNS resolvers. After identifying a censored domain, GFWatch can query the domain against popular DNS resolvers and verify whether their responses show any injection patterns.

GFWeb. GFWeb [87] is another system for performing measurements of HTTP and HTTPS censorship made by the GFW. Its goal is to augment existing measurements by including more domains and performing repeated trials over a longer time frame. Over a time span of 20 months, GFWeb scanned more than 1 billion domains and detected over 940K and over 50K censored domains for HTTP and HTTPS, respectively. The measurement approach designed by the authors allowed them to trigger HTTP(S) censorship without performing an actual three-way handshake. To this aim, they performed tests by sending SYN and PSH/ACK packet pairs, with the latter containing a HTTP GET or a TLS client hello. Additionally, GFWeb uses servers both inside and outside of China as vantage points, allowing it to scan both the ingress and egress behavior of the censor. This analysis demonstrated that the GFW acts in an asymmetrical manner, which was a previously unknown behavior. Other findings of this longitudinal study include changes in the filtering logic of the GFW. Previously there was considerable overblocking for some domains, which was found to be patched in later scans.

ICLab. This Internet monitoring platform was introduced by Niaki *et al.* [140]. The project measures connectivity to sites. This includes a test of the DNS request-response behavior with local and public resolvers, TCP connection test, certificate chain tests (for HTTPS URLs), an evaluation of the HTTP response (headers and body), a traceroute log, and a capture of all packets exchanged during the measurement procedure [140]. ICLab uses hosts located within the region that is being monitored as vantage points. These vantage points are split into 264 VPN endpoints and 17 endpoints operated by 17 volunteers. In both cases, these hosts run ICLab's measurement software (VPN points route traffic through a VPN while the volunteer-operated devices send the plain traffic probes) [140]. Since measurement results are compared by the authors, additional *control nodes*, located in the US, Europe, and Asia are employed. The data collection stopped in 2020, with traceroute data being updated until 2021.

IODA. The *Internet Outage Detection and Analysis* project [95] is a system designed for continuous monitoring of

the Internet. It is operated by the Internet Intelligence Lab at Georgia Tech's College of Computing. It detects large-scale outages at the network edge, including disruptions impacting ASes or entire regions within a country. IODA works in partnership with OONI [147]. It uses three distinct and complementary sources of Internet measurement data (BGP, Active Probing, and Internet Pollution) to detect outages and has provided near-real-time visualizations of Internet connectivity on a public platform. Regarding BGP, IODA monitors the reachability of IP addresses by analyzing routing data from global sources such as RouteViews and RIPE RIS [32]. It examines small network blocks to accurately detect whether parts of a network are offline or inaccessible. For active probing, IODA uses ICMP echo requests as part of the Trinocular method [162, 95]. When initial probes receive no response, the system sends additional probes to verify the network's status. The Trinocular methodology can accurately distinguish between online, offline, and uncertain states. Considering Internet pollution, IODA analyzes Internet Background Radiation (IBR) [32], which is unsolicited traffic resulting from misconfigurations, malware propagation, and network scanning. It filters out spoofed and bursty traffic (e.g., botnet scans) to extract a liveness signal based on distinct source IP addresses, enabling the detection of whether networks or regions are online or offline.

An example of the use of IODA is the investigation of complete Internet outages in Myanmar by Padmanabhan *et al.* in [153]. They analyzed the Coup Day in 2021. The outage was particularly evident in the BGP data sources, where the number of /24 address blocks reachable in Myanmar via BGP dropped from 695 to 376 - a 46 percent decrease. In [146] IODA data highlights Internet outages in Iran between September 2022 and October 2022. The BGP data source reflects routing announcements or global routing information exchanged between ISP routers, which are typically highly stable. A sudden drop, which in this case occurred frequently, indicates a disruption in Internet connectivity.

Kentik. Kentik is a large network observability company and offers a cloud-based NetFlow analytics platform designed for network monitoring, capable of analyzing traffic data collected from sources such as company network routers [146]. For example, Kentik provides a comprehensive overview of how data flows (traffic) within a network by leveraging aggregated NetFlow data from its customers.

NetFlow is a network protocol used by routers and switches to collect information about network traffic. It records metadata about IP traffic flows that pass through a NetFlow-enabled device (such as a router, switch, or host). The data collected includes: source and destination IP addresses of the data packets, protocols used (e.g., TCP, UDP), packet count, data volume in bytes, and connection duration [102]. This information can be aggregated to identify trends and patterns [43, 153]. The protocol and its logs provide detailed insights into network traffic,

supporting monitoring, optimization, and security. When NetFlow data is sourced from multiple locations, it offers a composite view of the Internet from various perspectives. Using this aggregated data, Kentik enables detailed analyses to draw conclusions about outages and other network disruptions. Kentik's customers, including tier-1 ISPs and global content providers, provide NetFlow data from Internet routers, making it possible to study censorship events and more [102]. As an example, Padmanabhan *et al.* used Kentik in [153] to analyze cellular data. To demonstrate that cellular connectivity had been heavily restricted, they compared traffic data from four major cellular providers with that of a leading non-cellular provider. Their analysis revealed a substantial reduction in traffic from cellular providers during both day and night. In contrast, the non-cellular provider experienced traffic drops only during the night. Further, in [146] Kentik was used to show that traffic volume from Facebook's network dropped substantially on September 21st, 2022, in Iran. On that day, the service WhatsApp, owned by Facebook, started showing signs of blocking.

The combination of IODA and Kentik offers significant advantages for detecting and analyzing network outages and censorship events. IODA specializes in macro-level analysis, monitoring unusual routing changes and identifying large-scale network outages across entire regions or ASes. It examines routing information to detect sudden changes in the reachability of network blocks. In contrast, Kentik focuses on micro-level analysis, providing granular insights into specific data patterns. This allows for detailed comparisons of data streams, complementing IODA's broader perspective. Together, these tools deliver a comprehensive approach to monitoring and analyzing network disruptions [146, 153].

OONI. The *Open Observatory of Network Interference* [71, 150] provides a *Measurement Aggregation Toolkit* (MAT) [148] that allows users to explore current and historic measurements of Internet censorship. MAT allows to investigate censorship specifically for selected ASes, countries, domains, dates and test methods. For instance, general Internet connectivity tests are available, but also tests for connecting to messenger services such as Facebook, Telegram, WhatsApp, or Signal. Further, certain middlebox-related tests, performance tests, tests for circumvention tools (Tor, including the Snowflake [26] pluggable transport, and Psiphon) as well as experimental tests (VPN connectivity and DNS check) can be analyzed regarding their results. OONI measurements are performed by volunteering users running *OONI Probe* [149] in a mobile, command line or desktop version. Further, a browser-based probe tool and an API to access OONI data are provided by the project. In general, differences between current and control measurements are considered as *anomalous* and can thus refer to a filtering attempt (invalid measurements can occur as well if a *measurement fails to execute or is unavailable for a particular query*) [48]. The

OONI website further provides timely reports on identified censorship events. Notably, a recent project has integrated OONI data into the Internet Yellow Pages (IYP), which is a knowledge graph with a unified access methodology [117]. The IYP allows an easier analysis of OONI data, as users do not need to autonomously perform any further data integration steps.

Psiphon. While Psiphon itself is an Internet censorship circumvention tool, the organization behind it also provides the *Psiphon Data Engine* (PDE) [161]. PDE is a platform that allows researchers to collectively analyze aggregated Internet censorship data. The platform collects censorship metrics, network efficiency measurements, connectivity tests and also collects records of social and political indicators regarding censorship events. So called “partners” can provide their measurement data to the platform. PDE provides information about the usage statistics of Psiphon like the unique users per day per country or the tunnel median round trip time. Currently, PDE lists OONI and M-Lab as partners. In consequence, we do not separately list PDE in Tab. 5.

RIPE Atlas. RIPE Atlas [170] is an Internet measurement platform allowing users to perform various connectivity measurements (e.g., ping, DNS, and traceroute) from over 13,000 active community-hosted probes. Despite not being directly linked with Internet censorship goals, RIPE ATLAS offers a diverse set of probes, which turned out to be valuable for implementing censorship measurements. For example, Anderson *et al.* [10] used RIPE Atlas to analyze blocking events in Russia and Turkey. Also, Gegenhuber *et al.* [76] used Atlas to measure deanonymization risks of tor users in Russia.

Tor Dataset. The Tor project publishes various usage data of the Tor network [186]. These include statistics about user numbers connecting to the network. These numbers can be used to detect censorship events. The Tor project itself uses this data to present the “Top-10 countries by possible censorship events”.

Triplet Censors. Niaki *et al.* [13] analyzed the DNS injection techniques of the GFW by observing the DNS results of the Alexa top 1M domains over a duration of nine months. They collected over 120 million forged DNS responses in their dataset. They published the dataset together with the code to allow for reproducibility studies and to foster the research in the domain.

Comparison of Platforms/Datasets. Tab. 5 compares the platforms and datasets discussed above. As can be seen, most of the datasets contain only DNS and HTTP data. Other probes are less frequent. In contrast, OONI provides a broad dataset that is driven by community measurements. Similarly, some datasets highlight only the censorship of a single country while others try to encompass

as many countries as possible, many of those rely on volunteers which has a direct impact on the number of probes. Most of the discussed datasets are still ongoing efforts, sometimes even providing real-time data to researchers. Only the Tor dataset is considered “passive” as it only contains usage statistics and not active probe traffic.

CensorLab. While not strictly a measurement platform or dataset, CensorLab [180] is a censorship experiment testbed. The goal of CensorLab is to enable researchers to emulate various censorship scenarios to evaluate existing, but also simulate and test newly developed censorship techniques, corresponding circumvention and measurement methods. This allows a proactive approach to censorship research instead of a reactive one. Other benefits of a simulated environment are reduced costs and reduced risks to measurement volunteers. CensorLab allows researchers to define the censors’ behavior through block lists and so called “censor programs” which can emulate a wide range of different censorship approaches.

Excluded Platforms and Datasets. We excluded a set of censorship platforms and datasets from our analysis due to superficial or dated measurement methodology or content, i.e., data older than 2015, or because they simply are not operated / existent anymore. For instance, in 2015, Aceto *et al.* [2] covered several measurement systems that we did not include in this paper, namely *ConceptDoppler*, *rTurtle*, *Herdict*, *Alkasir*, *YouTomb*, *Greatfire.org*, *CensMon*, *MOR*, *Weiboscope*, *Samizdat*, *UBICA*, *WCMT*, *Spookyscan* and *encore*. We further excluded the *OpenNet Initiative* (ONI) [152] that monitors Internet censorship in the sense of traffic filtering. Only a small fraction of data in the form of a summary CSV file is available anymore [151]. The latest update of the ONI dataset is dated Sep-2013 and was thus excluded from our analysis. However, an investigation of the dated dataset can be found in [78].

9. Trends and Future Challenges of Internet Censorship and its Measurement

In this section, we outline the major trends and future challenges of Internet censorship identified in the related literature. We highlight both technical and human aspects, considering changes in user behavior alongside the rapidly evolving landscape of software architectures and network protocols.

9.1. Technical Aspects

Measurement Design and Sampling. Quantifying the impact of censorship requires overcoming several challenges. The main issues, discussed a decade ago by Aceto and Pescapé [2], remain valid today. For instance, measurements should be characterized by an adequate probing frequency and must include a set of hosts/endpoints representative of a geographical region.

Table 5: Comparison platforms/datasets used for censorship measurement. (Abbreviations: TP: traffic probe, WCA: web-content analysis; BLK: blocked page analysis; Notes: ⁽¹⁾ 13 countries with volunteer-operated devices)

Dataset	CensoredPlanet	Encore	GFWatch	GFWeb	ICLab	IODA	Kentik	OONI	RIPE Atlas	Tor	Triplet Censors
Key Refs.	[38, 183]	[30]	[77, 89]	[87]	[140, 34]	[95]	[102]	[71]	[170]	[186]	[13]
Approach	active	active	active	active	active	active	active	active	active	passive	active
Method	TP,WCA,BLK	WCA,BLK	TP	TP	TP,BLK	TP,BLK	TP,BLK	TP,WCA,BLK	TP	-	TP,BLK
Latest Year	ongoing	2015	ongoing	ongoing	2021	ongoing	ongoing	ongoing	ongoing	ongoing	2020
# Countries	worldwide	170	1	1	62 ⁽¹⁾	worldwide	worldwide	worldwide	worldwide	worldwide	1
Probe Nodes	95,000+	88,260	2+	2+	281	500+	n/a	community based	community based 13,000+	n/a	1
HTTP(S) Probes	✓	✓	-	✓	✓	-	✓	✓	✓	-	-
QUIC Probes	-	-	-	-	-	-	-	✓([60])	-	-	-
DNS Probes	✓	✓	✓	-	✓	-	✓	-	✓	-	✓
DoT/DoH Probes	-	-	-	-	-	-	✓	✓	-	-	-
Tor Probes	-	-	-	-	-	-	-	✓	-	-	-
Tor Snowflake/WebRTC	-	-	-	-	-	-	-	✓	-	-	-
VPN Probes	-	-	-	-	✓	-	-	✓	-	-	-
Traceroute	✓	-	-	-	✓	only routing problems	✓	✓	✓	-	-

Collecting network traffic and behavior patterns is no longer primarily a problem of volume. In fact, advancements in storage and processing, often offered through as-a-Service models, have made it easier to obtain effective traffic snapshots compared to ten years ago.

Infrastructure, Routing, and AI Challenges. More pressing challenges stem from the evolving nature of the Internet architecture. First, there is a potential censorship inconsistency problem, where different ASes, ISPs or regions of the same country differ in their censorship policy implementation. As reported by Tsai *et al.* [189, 207], there is no global-scale ground truth available that can be used to validate a censorship dataset. Hence, certain authors and platforms (e.g., above-mentioned approach of CenDetect) explicitly address how they handle false-positives and false-negatives. Second, routing path inconsistencies complicate traffic probing, often due to the ubiquitous diffusion of CDNs and load balancing components. As a partial workaround, advancements in AI to process large-scale complex information may support researchers in Internet censorship, e.g., to “repair” traffic plagued by mismatches. Yet, some state-level censors are expected to deploy mechanisms to influence the interaction with AI itself. For instance, censors may operate over the communication layer, e.g., by blocking certain IP ranges to prevent the availability of complete datasets for training Large Language Models (LLMs) or block access to these platforms, such as in China, Russia, Iran, and Israel [19, 72]. Therefore, future research should consider the investigation of LLM-based censorship (as, e.g., initially done by Ahmed and Knockel [3]). Considering LLMs (and AI in general), another expected trend is that censors work towards applying AI-based detection tools for nuanced critiques in textual and/or visual form [40], for which research is still in its infancy. Similar advancements can be anticipated for audio/video transmissions and content.

Evolving Targets and Services. Another important challenge to face in the near future entails the rapid evolution of targets and services. Even if some works have shown

how to promptly adapt to new communication protocols (e.g., for QUIC [60]), the actual panorama of Internet services is changing at an unprecedented pace. Social media platforms, including gaming and video services, have tight popularity cycles, thus demanding for continuous modifications of the measurement techniques/algorithms.

In addition, measurement properties can shift due to alternating user-bases or methodological adjustments within datasets [48].

Integrating and Combining Datasets. To address these challenges, researchers conducted dataset-overlapping work. Crowder *et al.* [48] incorporated multiple of the above-mentioned datasets (Satellite, Hyperquack HTTP(S), Tor, OONI, and GFWatch) into a framework for joint statistical analysis to investigate censorship events. Similarly, Lipphardt *et al.* [117] integrated OONI data into the IYP and mention that they plan to integrate data from additional measurement platforms in order to provide a unified accessibility through the query language *Cypher*.

Adversarial Dynamics and Arms Race. The relevant corpus of works and the vivacity of censorship circumvention research is expected to fuel an “arms race” among states, regimes, researchers, and activists. In this vein, it is important to underline the suggestion of Aceto and Pescapé reported in [2] stating that measurement approaches should consider that censors (could) identify probing devices and operate to actively hide the effects of blockages or filterings. This adversarial dynamic is likely to intensify as measurement and circumvention techniques evolve.

Research Quality and Education. A final remark to improve the technical research is to plan for proper strategies to ensure quality and consistency. First, experiments devoted to gather Internet traffic should follow clear quality criteria [156]. Gained results should then follow the *FAIR* principle, i.e., research work (including, code, data, metadata, notebooks etc.) should be *F*indable, *A*ccessible, *I*nteroperable, and *R*eusable. A way to aid re-usability is

to publish the artifacts on open repositories and provide companion software or traffic capture metadata.

Finally, we encourage starting to teach Internet censorship (measurement and circumvention) as a complement in network security and privacy curricula, especially to raise awareness and reduce inconsistent terminologies and scientific re-inventions [200].

9.2. Societal and Human Aspects

Beyond the Network Layer. So far, we almost exclusively covered network-level techniques. However, to provide a deeper understanding of the topic, it is beneficial to also understand societal and human aspects of Internet censorship. While, technical measurements map the prevalence of restrictions across different regions, the success, and accuracy also depends on human aspects involved. This includes how people adapt their online behavior in response to perceived threats, trust, fear, and resistance.

Although digital technologies have transformed public discourse and how communities connect, the Internet is not a neutral communication tool. More so, it is a structure shaped by cultural differences, power imbalances and economic dynamics [33, 198, 214]. These factors influence how the Internet is constructed, controlled and experienced, with varying levels of censorship creating tensions between national regulations and its originally borderless nature [198]. Additionally, geographical factors, such as historical legacy, legal frameworks, governmental structures, and societal norms, shape distinct patterns of restriction, surveillance, and control [198]. In turn, the shift of power dynamics from nation states to global networks further emphasizes the role of human behavior in the success of censorship techniques, measurement, and circumvention practices.

Motivations for Censorship Shape Measurement Challenges. Censorship is justified in diverse ways, shaping both enforcement mechanisms and user responses. Motivations can be broadly categorized into four main categories [2, 128, 93, 11]: (i) *Political Control*, (ii) *Social and Moral Regulation*, (iii) *Parental and Company Controls*, and (iv) *Economic Motivations*. Each motive influences what is censored, how it is applied, and how it is perceived. Political and social motives often fluctuate around events such as elections or crises, while economic and parental controls create subtler but still significant restrictions.

(i) *Political Control:* Governments, particularly repressive ones, strategically control Internet access to unwanted political opinions (e.g., NGOs, websites of foreign newspapers or social media). The aim is to limit its potential to empower citizens and prevent them from bypassing state-controlled media [90, 173, 50, 7, 14]. Motivations for censorship range from suppressing political dissent and human rights activism to silencing criticism of the state or its officials, as seen in countries like China and Iran [106, 85].

Social media platforms, vital spaces for political mobilization, are often restricted or shut down entirely [6, 79, 159]. In such cases, measuring censorship requires more than merely detecting blocked news websites, but also involves understanding how fear influences open discourse and indirect forms of online manipulation [75, 143]. This is especially crucial from a temporal perspective, as censorship practices strongly fluctuate around sensitive events like elections or protests [72].

(ii) *Social and Moral Regulation:* Another governmental justification for censorship is expressing social and moral concerns. By claiming to protect societal values, public morality, or vulnerable groups [106, 111, 119], mostly illegal content is targeted. This includes blocking websites linked to illicit trade (e.g., drugs, weapons), child exploitation, pornography, or platforms promoting hate speech. The definition of harmful content or what is deemed immoral varies across cultures [105, 85], leading to uneven enforcement that disproportionately impacts marginalized groups [198, 45, 31]. Beyond safeguarding public morality, concerns about terrorism, national security, preventing crimes and maintaining social stability are increasingly cited [196, 50, 130]. However, the relationship between national security laws and the Internet is complex, encompassing measures such as content takedowns and network shutdowns, which are responses to local conditions, such as armed conflicts or political unrest, reflecting historical tensions and contemporary power struggles. The use of centralized power structures to monitor and shape online content to very different extents, in some cases, is therefore rooted in fear of losing legitimacy [50].

(iii) *Parental and Company Controls:* Parental controls are often implemented at the household level or through private software filters, but some governments mandate or encourage Internet service providers or software vendors to offer content controls for minors [172, 114, 54]. Although less intrusive than broader state censorship, these measures align with local regulatory environments.

(iv) *Economic Motivations:* Other motivations for censorship can involve intellectual property protections, such as blocking illegal downloads. Companies and governments may also block competitor websites or restrict traffic to protect trade secrets [129]. Although this practice can serve legitimate security purposes, it can also become a rationale to censor or restrict broader sets of content [163, 52]. Such overblocking of potential legitimate content could affect users' trust in information sources [52, 27]. Regions with strong government-industry ties or heightened espionage concerns tend to enforce such restrictions more aggressively, reflecting the influence of economic and political alliances on censorship patterns [131, 192]. Additionally, enforcing censorship often leads to unintended non-technical consequences, such as stifled innovation.

Managing Public Attitudes as a Form of Control. Beyond technical restrictions, governments actively shape public opinions to legitimize censorship. For example, state-controlled Chinese media frequently portray the Internet in a negative light, reinforcing public acceptance of Internet restrictions [63, 196]. This reflects the concept of *networked authoritarianism* [120], where authoritarian governments allow some degree of connectivity but strategically manage and tightly control online discourse to reinforce state narratives. The example of China demonstrates how deeply rooted cultural norms can lead citizens to regard the government as the sole legitimate authority, thereby making its actions appear more natural and expected [63, 28]. In this context, governments also refer to user studies, stating that the majority of respondents are in favor of Internet censorship and support government control [94, 81, 196]. However, this high acceptance rate of Internet censorship is irrespective of the nuances stated by respondents about which types of content should be censored (e.g., pornographic content over online chatting), [196, 81]. Nonetheless, the study by Wang and Mark [196] also showed that over time, even if there is an uprising against censorship at first, normalization reduces resistance, making censorship more effective even without broad technical blocking. Consequently, individuals who adapt and modify their online behavior, impact censorship detection and measurement, as such subtle changes leave fewer visible traces.

User Behavior and Self-Censorship. As already mentioned, whereas self-censorship is primarily an outcome of censorship and surveillance rather than a deliberate technique by censors, one psychological motivator that reinforces it is fear of isolation: the concern that expressing an unpopular opinion may lead to social exclusion or disapproval. This fear can lead to a heightened tendency to withhold opinions, reflecting the willingness to self-censor [2, 29] (see Sect. 4). Moreover, self-censorship is closely related to the *chilling effect*, where users adapt their online behavior, and the *extended chilling effect*, where such adaptations spill over into offline life [215, 121]. Beyond these individual-level processes, the *spiral of silence* [144, 74] provides a broader social explanation: users are further discouraged from engaging in discussions when they perceive their opinions to be in the minority, even in the absence of direct censorship [91]. As more people stay silent, the dominant view appears even stronger, which in turn pressures more individuals into self-censorship, reinforcing the spiral of silence and making certain views disappear from public discourse [84, 29, 74]. Such social conformity amplifies the chilling effect, alters Internet usage, and can mislead passive censorship detection systems that ignore behavioral adaptations [2]. Consequently, technical measurements may underestimate censorship levels when they fail to account for human behavior.

Challenges in Data Collection and Interpretations. Collecting human-centered data introduces biases as well as

ethical and methodological risks. For example, OONI probe usage spikes during censorship events, meaning anomalies may correlate with participant activity rather than censorship prevalence: *users run OONI more often when they are experiencing censorship* [48]. Simultaneously, the use of Volunteer-Operated Devices (VODs) raises *ethical and logistical* challenges [140]. In highly censored environments, participating in data gathering, such as contributing to measurement platforms or reporting blocked content, can invite legal reprisal, political or social risks [52]. This produces self-selection bias, since only certain groups contribute [177, 196]. Marwick and Boyd [123] found that people who have been systematically and structurally marginalized (e.g., LGBTQ communities, people of color) are particularly sensitive to perceived surveillance, leading to heightened privacy concerns and behavioral adaptation, such as underreporting [158]. Thus, researchers bear responsibility to protect volunteers, assess political risks and work with trusted organizations [140]. Another challenge, also influencing potential *data misinterpretation*, arises when unskilled users are unable to distinguish between malfunction and censorship. As pointed out by Aceto *et al.* [2], users might blame a destination host for not delivering a website if no blocking page is shown, although a censor could slow down or block a connection to a particular target. Similarly, users may assume their Internet service provider or destination system is at fault, and thus censorship might be attributed to an end-user's client device. Additionally, some forms of content restriction, such as algorithmic filtering or social media shadowbans, may not be easily detectable. Therefore, to avoid missing the nuanced reasons behind traffic anomalies and to ensure correct data interpretation, local expertise and contextual awareness should be considered [183, 104]. Overall, while technical approaches effectively detect direct censorship, they risk underestimating suppression when human behavior and societal factors are ignored. Addressing these challenges by integrating technical, psychological and societal perspectives will not only improve measurement accuracy, but also support those most affected by restrictions.

10. Conclusion

We surveyed the capabilities of Internet censors. To this end, we introduced a generic taxonomy of censorship systems and provided an in-depth coverage of censorship methodology on the basis of network protocols. We have shown how current censorship is conducted, incorporating sophisticated methods targeting BGP, TLS, and several forms of circumvention tools. We provided a detailed explanation of censorship *measurement* on the Internet, transport, and application layers, explained how the measurement of circumvention tool censorship can be conducted and compared available censorship datasets and platforms. Finally, we discussed trends and challenges, including human aspects.

Overall, several censoring and measurement methods of the Internet, transport and application layers remained consistent for years. However, more of the community gained a better understanding of measurement limitations and developed more sophisticated methods that aim to improve precision and accuracy of experiments. Recent techniques incorporated novel network protocols, platforms and technological trends. Another major novelty of the last years is the focus on societal and human aspects (both, from a censorship and from a measurement perspective), which has been a side issue in the past.

Declaration of interests

The authors 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. We like to thank the students of Ulm University's class on Internet Censorship for the vivid discussions and the feedback gained that influenced parts of this paper. We also like to thank Paul Prechtel (Ulm Univ.) for his feedback on the BGP sections of this paper. Finally, we would like to express our thankfulness for the reviewers' feedback, which led to the improvement of the original manuscript.

References

- [1] Aase, N., Crandall, J.R., Díaz, Á., Knockel, J., Molinero, J.O., Saia, J., Wallach, D.S., Zhu, T., 2012. Whiskey, Weed, and Wukan on the World Wide Web: On measuring censors' resources and motivations, in: Proc. FOCI.
- [2] Aceto, G., Pescapé, A., 2015. Internet censorship detection: A survey. *Computer Networks* 83, 381–421. URL: <https://www.sciencedirect.com/science/article/pii/S1389128615000948>, doi:<https://doi.org/10.1016/j.comnet.2015.03.008>.
- [3] Ahmed, M., Knockel, J., 2024. The impact of online censorship on LLMs. Proc. FOCI '24.
- [4] Al-Musawi, B., Branch, P., Armitage, G., 2017. BGP anomaly detection techniques: A survey. *IEEE Communications Surveys & Tutorials* 19, 377–396. URL: <http://ieeexplore.ieee.org/document/7723902/>, doi:[10.1109/COMST.2016.2622240](https://doi.org/10.1109/COMST.2016.2622240).
- [5] Al-Musawi, B., Shehada, D., Hussain, A.J., 2025. Novel data-driven geolocation approach for detecting smuggled internet traffic. *IEEE Access* 13, 36306–36320. URL: <https://ieeexplore.ieee.org/document/10896665/?arnumber=10896665>, doi:[10.1109/ACCESS.2025.3543876](https://doi.org/10.1109/ACCESS.2025.3543876).
- [6] Al-Zaman, Sayeed, M., 2024. Do the reasons for social media censorship vary across political regimes and social media platforms? *Communication Reports*, 1–14doi:[10.1080/08934215.2024.2432320](https://doi.org/10.1080/08934215.2024.2432320).
- [7] Al-Zaman, M.S., Shiblee Noman, M.M., 2024. Rise of digital authoritarianism? exploring global motivations behind governmental social media censorship. *Javnost-The Public* 31, 529–544.
- [8] Amich, A., Eshete, B., Yegneswaran, V., Hoang, N.P., 2023. DeResistor: Toward detection-resistant probing for evasion of Internet censorship, in: Proc. USENIX Security Symposium 2023, pp. 2617–2633. URL: <https://www.usenix.org/conference/usenixsecurity23/presentation/amich>.
- [9] Anderson, C., 2013. Dimming the Internet: Detecting throttling as a mechanism of censorship in iran. URL: <http://arxiv.org/abs/1306.4361>, arXiv:1306.4361.
- [10] Anderson, C., Winter, P., Roya, 2014. Global network interference detection over the RIPE atlas network, in: 4th USENIX Workshop on Free and Open Communications on the Internet (FOCI 14), USENIX Association, San Diego, CA. URL: <https://www.usenix.org/conference/foci14/workshop-program/presentation/anderson>.
- [11] Ang, P.H., Nadarajan, B., 1996. Censorship and the Internet: A Singapore perspective. *Comm. ACM* 39, 72–78.
- [12] Anonymous, 2014. Towards a comprehensive picture of the great firewall's DNS censorship, in: FOCI '14.
- [13] Anonymous, Niaki, A.A., Hoang, N.P., Gill, P., Houmansadr, A., 2020. Triplet censors: Demystifying great firewall's DNS censorship behavior, in: Proc. FOCI '20.
- [14] Ashokkumar, A., Talaifar, S., Fraser, W., Landabur, R., Buhrmester, M., Gómez, A., Paredes, B., Swann, W., 2020. Censoring political opposition online: Who does it and why. *Journal of Experimental Social Psychology* 91, 104031. doi:[10.1016/j.jesp.2020.104031](https://doi.org/10.1016/j.jesp.2020.104031).
- [15] Ballani, H., Francis, P., Zhang, X., 2007a. A study of prefix hijacking and interception in the Internet. *ACM SIGCOMM Computer Communication Review* 37, 265–276.
- [16] Ballani, H., Francis, P., Zhang, X., 2007b. A study of prefix hijacking and interception in the Internet. *ACM SIGCOMM Computer Communication Review* 37, 265–276. URL: <https://dl.acm.org/doi/10.1145/1282427.1282411>, doi:[10.1145/1282427.1282411](https://doi.org/10.1145/1282427.1282411).
- [17] Barnes, R., Bhargavan, K., Lipp, B., Wood, C.A., 2022. Hybrid Public Key Encryption. RFC 9180. URL: <https://www.rfc-editor.org/info/rfc9180>, doi:[10.17487/RFC9180](https://doi.org/10.17487/RFC9180).
- [18] Becker, K.B., 2011. Internetzensur in China. VS Verlag für Sozialwissenschaften. In German.
- [19] Berger, H., Shavitt, Y., 2024. Measuring DNS censorship of generative AI platforms. URL: <https://arxiv.org/abs/2412.14286>, arXiv:2412.14286.
- [20] Bhaskar, A., Pearce, P., 2022. Many roads lead to Rome: How packet headers influence DNS censorship measurement, in: Proc. USENIX Security '22.
- [21] Bhaskar, A., Pearce, P., 2024. Understanding routing-induced censorship changes globally, in: Proc. 31st ACM Conf. Computer and Communications Security (CCS'24), ACM.
- [22] Bishop, M., 2019. Computer Security. Art and Science. 2 ed., Pearson Education.
- [23] Bitso, C., Fourie, I., Bothma, T.J., 2013. Trends in transition from classical censorship to Internet censorship: selected country overviews. *Innovation: journal of appropriate librarianship and information work in Southern Africa* 2013, 166–191.
- [24] Bock, K., Bharadwaj, P., Singh, J., Levin, D., 2021. Your censor is my censor: Weaponizing censorship infrastructure for availability attacks, in: 2021 IEEE Security and Privacy Workshops (SPW), pp. 398–409. doi:[10.1109/SPW53761.2021.00059](https://doi.org/10.1109/SPW53761.2021.00059).
- [25] Bock, K., Hughey, G., Qiang, X., Levin, D., 2019. Geneva: Evolving censorship evasion strategies, in: Proc. 2019 ACM SIGSAC Conf. Computer and Communications Security, ACM. p. 2199–2214. URL: <https://doi.org/10.1145/3319535.3363189>, doi:[10.1145/3319535.3363189](https://doi.org/10.1145/3319535.3363189).
- [26] Bocovich, C., Breault, A., Fifield, D., Wang, X., et al., 2024. Snowflake, a censorship circumvention system using temporary WebRTC proxies, in: Proc. USENIX Security '24, pp. 2635–2652.
- [27] Brown, I., 2008. Internet censorship: be careful what you ask for. Bahcesehir University Press. pp. 74–91.
- [28] Burgers, T., Robinson, D.R.S., 2016. Networked authoritarianism is on the rise. *Sicherheit und Frieden (S+F) / Security and Peace* 34, 248–252. URL: <http://www.jstor.org/stable/26429018>.
- [29] Burnett, A., Knighton, D., Wilson, C., 2022. The self-censoring majority: How political identity and ideology impacts willingness to self-censor and fear of isolation in the

- United States. *Social Media + Society* 8. doi:10.1177/20563051221123031.
- [30] Burnett, S., Feamster, N., 2015. Encore: Lightweight measurement of web censorship with cross-origin requests, in: *Proc. ACM Conf. Special Interest Group on Data Communication*, ACM. p. 653–667.
- [31] Busch, A., Theiner, P., Breindl, Y., 2018. Internet censorship in liberal democracies: Learning from autocracies? Managing democracy in the digital age: Internet regulation, social media use, and online civic engagement, 11–28.
- [32] CAIDA Project, 2024. Website. URL: <https://www.caida.org/projects/ioda/>.
- [33] Cairncross, F., 2002. The death of distance. *RSA Journal* 149, 40–42. URL: <http://www.jstor.org/stable/41380436>.
- [34] Calipr Networking Group, 2025. Information controls lab. <https://iclab.gitlab.io>.
- [35] Calle, P., Savitsky, L., Bhagoji, A.N., Hoang, N.P., Cho, S., 2024. Toward automated DNS tampering detection using machine learning, in: *Proc. FOCI '24*.
- [36] Castro, R., Coates, M., Liang, G., Nowak, R., Yu, B., 2004. Network tomography: Recent developments. *Statistical Science* 19, 499–517. doi:10.1214/088342304000000422.
- [37] Caviglione, L., 2021. Trends and challenges in network covert channels countermeasures. *Applied Sciences* 11, 1641.
- [38] Censored Planet Team, 2023. Censored Planet: An internet-wide, longitudinal censorship observatory. <https://censoredplanet.org/>. Accessed: February 12, 2025.
- [39] Chai, Z., Ghafari, A., Houmansadr, A., 2019. On the importance of Encrypted-SNI (ESNI) to censorship circumvention, in: *Proc. FOCI '19*. URL: <https://www.usenix.org/conference/foci19/presentation/chai>.
- [40] Chen, Y., . The Accuracy and Biases of AI-Based Internet Censorship in China 4, 27–36. URL: <https://www.pioneerpublisher.com/jrssh/article/view/1223>.
- [41] Cherubin, G., Jansen, R., Troncoso, C., 2022. Online website fingerprinting: Evaluating website fingerprinting attacks on tor in the real world, in: *Proc. USENIX Security '22*, pp. 753–770.
- [42] Cho, S., Fontugne, R., Cho, K., Dainotti, A., Gill, P., 2019. BGP hijacking classification, in: *2019 Network Traffic Measurement and Analysis Conf. (TMA)*, pp. 25–32.
- [43] Cisco Systems, 2024. Configuring NetFlow and NetFlow data export. URL: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/netflow/configuration/xr-16/test/nf-xe-16-book/cfg-nflow-data-expt-xe.html?utm_source=chatgpt.com.
- [44] Citizen Lab, et al., 2014. URL testing lists intended for discovering website censorship. URL: <https://github.com/citizenlab/test-lists>. <https://github.com/citizenlab/test-lists>.
- [45] Clark, J.D., Faris, R.M., Morrison-Westphal, R.J., Noman, H., Tilton, C.B., Zittrain, J.L., 2017. The shifting landscape of global Internet censorship. *Berkman Klein Center for Internet & Society Scholarly Articles* [105] .
- [46] Clayton, R., Murdoch, S.J., Watson, R.N.M., 2006. Ignoring the great firewall of China, in: *Privacy Enhancing Technologies*, Springer. pp. 20–35.
- [47] Collins, M., 2017. *Network Security Through Data Analysis*. 2 ed., O'Reilly.
- [48] Crowder, A., Olszewski, D., Traynor, P., Butler, K.R.B., 2024. I can show you the world (of censorship): Extracting insights from censorship measurement data using statistical techniques, in: *Annual Computer Security Applications Conf. (ACSAC'24)*, IEEE. pp. 1123–1138. doi:10.1109/ACSAC63791.2024.00091.
- [49] Dainotti, A., Squarcella, C., Aben, E., Claffy, K.C., Chiesa, M., Russo, M., Pescapé, A., 2011. Analysis of country-wide Internet outages caused by censorship, in: *Proc. Internet Measurement Conference (IMC '11)*, ACM. doi:10.1145/2068816.2068818.
- [50] Dal, A., Nisbet, E.C., 2022. Walking through firewalls: Circumventing censorship of social media and online content in a networked authoritarian context. *Social Media + Society* 8.
- [51] Davitt, K., Ristea, D., Russell, D., Murdoch, S.J., 2024. CoStricTor: Collaborative HTTP Strict Transport Security in Tor Browser. *Proc. Privacy Enhancing Technologies* 2024, 343–356.
- [52] Deibert, R., Palfrey, J., Rohozinski, R., Zittrain, J.L., 2008. *Access Denied: The Practice and Policy of Global Internet Filtering*. MIT Press. doi:10.7551/mitpress/7617.001.0001.
- [53] Detal, G., Hesmans, B., Bonaventure, O., Vanaubel, Y., Donnet, B., 2013. Revealing middlebox interference with tracebox, in: *Proc. 2013 Conf. Internet Measurement Conference*, ACM. pp. 1–8. URL: <https://dl.acm.org/doi/10.1145/2504730.2504757>, doi:10.1145/2504730.2504757.
- [54] Die Bundesländer, 2024. Sechster Staatsvertrag zur Änderung medienrechtlicher Staatsverträge (Entwurf). <https://www.ministerpraesident.sachsen.de/ministerpraesident/TOPI-10-Sechster-Medienänderungsstaatsvertrag.pdf>.
- [55] Dingledine, R., Mathewson, N., Syverson, P.F., et al., 2004. Tor: The second-generation onion router, in: *Proc. USENIX Sec. Symp. '04*, pp. 303–320.
- [56] Dischinger, M., Mislove, A., Haeberlen, A., Gummadi, K.P., 2008. Detecting Bittorrent blocking, in: *Proc. Conf. Internet Measurement (IMC '08)*, ACM. doi:10.1145/1452520.1452523.
- [57] Dunna, A., O'Brien, C., Gill, P., . Analyzing China's Blocking of Unpublished Tor Bridges. URL: <https://www.usenix.org/conference/foci18/presentation/dunna>.
- [58] Eastlake 3rd, D.E., 2011. Transport Layer Security (TLS) Extensions: Extension Definitions. RFC 6066. URL: <https://www.rfc-editor.org/info/rfc6066>, doi:10.17487/RFC6066.
- [59] Ejeta, T.G., Kim, H.J., 2017. Website fingerprinting attack on psiphon and its forensic analysis, in: *Proc. International Workshop on Digital Forensics and Watermarking (IWDW)*, Springer. pp. 42–51.
- [60] Elmenhorst, K., Schütz, B., Aschenbruck, N., Basso, S., 2021. Web censorship measurements of HTTP/3 over QUIC, in: *Proc. 21st ACM Internet Measurement Conf. (IMC)*, ACM. p. 276–282. URL: <https://doi.org/10.1145/3487552.3487836>, doi:10.1145/3487552.3487836.
- [61] Ensafi, R., Knockel, J., Alexander, G., Crandall, J.R., 2014. Detecting intentional packet drops on the internet via TCP/IP side channels, in: *Passive and Active Measurement*, Springer International Publishing. pp. 109–118.
- [62] Ermoshina, K., Musiani, F., 2021. The Telegram ban: How censorship “made in Russia” faces a global Internet. First Monday URL: <https://firstmonday.org/ojs/index.php/fm/article/view/11704>, doi:10.5210/fm.v26i5.11704.
- [63] Fallows, D., 2008. Most Chinese say they approve of government Internet control. *Pew Internet & American Life Project*, Washington.
- [64] Fan, S., Sippe, J., San, S., Sheffey, J., Fifield, D., Houmansadr, A., Wedwards, E., Wustrow, E., 2025. Wallbleed: A Memory Disclosure Vulnerability in the Great Firewall of China, in: *Network and Distributed System Security (NDSS) Symposium*. doi:10.14722/ndss.2025.230237.
- [65] Feng, Y., Sion, R., 2023. A Study of China's Censorship and Its Evasion Through the Lens of Online Gaming, in: *USENIX Security Symposium 2023*.
- [66] Fifield, D., 2017. Threat modeling and circumvention of Internet censorship. Ph.D. thesis. Electrical Engineering and Computer Sciences, University of California at Berkeley.
- [67] Fifield, D., 2020. Turbo Tunnel, a good way to design censorship circumvention protocols, in: *Proc. FOCI '20*. URL: <https://www.usenix.org/conference/foci20/presentation/fifield>.
- [68] Fifield, D., Lan, C., Hynes, R., Wegmann, P., Paxson, V., 2015. Blocking-resistant communication through domain fronting. *Proc. Privacy Enhancing Technologies* .
- [69] Fifield, D., Tsai, L., 2016. Censors' delay in blocking circumvention proxies, in: *Proc. FOCI'16, USENIX*.
- [70] Figueira, G., Barradas, D., Santos, N., 2022. Stegozoa: Enhancing WebRTC covert channels with video steganography

- for Internet censorship circumvention, in: Proc. 2022 ACM on Asia Conf. Computer and Communications Security (AsiaCCS), pp. 1154–1167.
- [71] Filasto, A., Appelbaum, J., 2012. OONI: open observatory of network interference, in: Proc. FOCI '12.
- [72] Freedom House, 2024. Freedom of the net 2024. <https://freedomhouse.org/sites/default/files/2024-10/FREEDOM-ON-THE-NET-2024-DIGITAL-BOOKLET.pdf>.
- [73] Gao, P., Li, G., Shi, Y., Wang, Y., 2020. VPN traffic classification based on payload length sequence, in: 2020 International Conf. Networking and Network Applications (NaNA), pp. 241–247. URL: <https://ieeexplore.ieee.org/document/9353766/?arnumber=9353766>, doi:10.1109/NaNA51271.2020.00048.
- [74] Gearhart, S., Adegbola, O., 2025. Spiral of silence and social conflict. The Handbook of Social and Political Conflict, 141.
- [75] Gebhart, G., Kohno, T., 2017. Internet censorship in Thailand: User practices and potential threats, in: 2017 IEEE European symposium on security and privacy (EuroS&P), IEEE. pp. 417–432.
- [76] Gegenhuber, G.K., Maier, M., Holzbauer, F., Mayer, W., Merzdovnik, G., Weippl, E., Ullrich, J., 2023. An extended view on measuring tor as-level adversaries. Computers & Security 132, 103302. URL: <https://www.sciencedirect.com/science/article/pii/S0167404823002122>, doi:https://doi.org/10.1016/j.cose.2023.103302.
- [77] GFWatch, 2025. project website. URL: <https://gfwatch.org/>.
- [78] Gill, P., Crete-Nishihata, M., Dalek, J., Goldberg, S., Senft, A., Wiseman, G., 2015. Characterizing web censorship worldwide: Another look at the OpenNet initiative data. ACM Trans. Web (TWEB) 9, 1–29.
- [79] Gillett, R., Gray, J., Kaye, D.B., 2023. ‘just a little hack’: Investigating cultures of content moderation circumvention by facebook users. New Media & Society 26, 146144482211476. doi:10.1177/14614448221147661.
- [80] Guo, H., Heidemann, J., 2017. Detecting ICMP rate limiting in the Internet (extended). USC/Information Sciences Institute, Los Angeles, CA, Tech. Rep. ISI-TR-717.
- [81] Guo, S., Feng, G., 2012. Understanding support for internet censorship in China: An elaboration of the theory of reasoned action. J OF CHIN POLIT SCI 17, 33–52. doi:https://doi.org/10.1007/s11366-011-9177-8.
- [82] Hall, J.L., Aaron, M.D., Andersdotter, A., Jones, B., Feamster, N., Knodel, M., 2023. A survey of worldwide censorship techniques. RFC 9505. URL: <https://www.rfc-editor.org/info/rfc9505>, doi:10.17487/RFC9505.
- [83] Handley, M., Paxson, V., Kreibich, C., 2001. Network intrusion detection: Evasion, traffic normalization, and end-to-end protocol semantics, in: Proc. 10th USENIX Security Symposium, pp. 115–131. URL: https://www.usenix.org/legacy/events/sec01/full_papers/handley/handley.pdf.
- [84] Hayes, A.F., 2007. Exploring the forms of self-censorship: On the spiral of silence and the use of opinion expression avoidance strategies. Journal of Communication 57, 785–802. doi:https://doi.org/10.1111/j.1460-2466.2007.00368.x.
- [85] He, K., Eldridge, S., Broersma, M., 2024. Tactics of disconnection: How netizens navigate china’s censorship system. Media and Communication 12, 8670. doi:https://doi.org/10.17645/mac.8670.
- [86] Hoang, N., Niaki, A., Dalek, J., Knockel, J., Lin, P., Marczak, B., Crete-Nishihata, M., Gill, P., Polychronakis, M., 2021a. How great is the great firewall? measuring China’s DNS censorship, in: Proc. USENIX Security Symp. '21.
- [87] Hoang, N.P., Dalek, J., Crete-Nishihata, M., Christin, N., Yegneswaran, V., Polychronakis, M., Feamster, N., 2024. GFWeb: Measuring the Great Firewall’s web censorship at scale, in: 33rd USENIX Security Symposium (USENIX Security 24), USENIX Association. pp. 2617–2633. URL: <https://www.usenix.org/conference/usenixsecurity24/presentation/hoang>.
- [88] Hoang, N.P., Doreen, S., Polychronakis, M., 2019. Measuring I2P censorship at a global scale, in: Proc. FOCI '19.
- [89] Hoang, N.P., Niaki, A.A., Dalek, J., Knockel, J., Lin, P., Marczak, B., Crete-Nishihata, M., Gill, P., Polychronakis, M., 2021b. How great is the great firewall? measuring china’s DNS censorship, in: Proc. USENIX Security '21, pp. 3381–3398.
- [90] Hobbs, W.R., Roberts, M.E., 2018. How sudden censorship can increase access to information. American Political Science Review 112, 621–636. doi:10.1017/S0003055418000084.
- [91] Hoffmann, C.P., Lutz, C., 2017. Spiral of silence 2.0: Political self-censorship among young facebook users, in: Proc. 8th international Conf. social media & society, pp. 1–12.
- [92] Holz, R., Riedmaier, T., Kammenhuber, N., Carle, G., 2012. X.509 forensics: Detecting and localising the SSL/TLS Man-in-the-Middle, in: Computer Security – ESORICS 2012, Springer Berlin Heidelberg. pp. 217–234. doi:10.1007/978-3-642-33167-1_13.
- [93] Houmansadr, A., Brubaker, C., Shmatikov, V., 2013. The parrot is dead: Observing unobservable network communications, in: 2013 IEEE Symposium on Security and Privacy, IEEE. pp. 65–79.
- [94] Internet Society, 2012. Global Internet user survey 2012. URL: <https://www.internetsociety.org/internet/global-internet-user-survey-2012/>.
- [95] IODA Project, 2024. Website. URL: <https://ioda.inetintel.cc.gatech.edu/>.
- [96] Iszaevich, G.E.W., 2019. Distributed detection of Tor directory authorities censorship in Mexico, in: Proc. International Conf. Networks (ICN 2019), Iaria. pp. 82–86. URL: <https://ru.iiec.unam.mx/4538/>.
- [97] Jermyn, J., Weaver, N., 2017. Autosonda: Discovering Rules and Triggers of Censorship Devices, in: 7th USENIX Workshop on Free and Open Communications on the Internet (FOCI 17).
- [98] Jin, L., Hao, S., Wang, H., Cotton, C., 2021a. Understanding the impact of encrypted DNS on Internet censorship, in: Proc. Web Conf. (WWW'21), ACM. p. 484–495. URL: <https://doi.org/10.1145/3442381.3450084>, doi:10.1145/3442381.3450084.
- [99] Jin, L., Hao, S., Wang, H., Cotton, C., 2021b. Understanding the practices of global censorship through accurate, end-to-end measurements. Proc. ACM Meas. Anal. Comput. Syst. 5. URL: <https://doi.org/10.1145/3491055>, doi:10.1145/3491055.
- [100] Kakhki, A., Razaghpanah, A., Li, A., Koo, H., Golani, R., Choffnes, D., Gill, P., Mislove, A., 2015. Identifying traffic differentiation in mobile networks, in: Proc. 2015 Internet Measurement Conf. (IMC), ACM.
- [101] Karunanayake, I., Ahmed, N., Malaney, R., Islam, R., Jha, S.K., 2021. De-anonymisation attacks on Tor: A survey. IEEE Communications Surveys & Tutorials 23, 2324–2350.
- [102] Kentik Technologies, 2024. Website. URL: <https://kentic.com>.
- [103] Khattak, S., Elahi, T., Simon, L., Swanson, C.M., Murdoch, S.J., Goldberg, I., 2016. SoK: Making sense of censorship resistance systems. Proc. Privacy Enhancing Technol. 2016, 37–61. doi:10.1515/popets-2016-0028.
- [104] Khattak, S., Javed, M., Anderson, P.D., Paxson, V., 2013. Towards illuminating a censorship monitor’s model to facilitate evasion, in: Proc. FOCI '13.
- [105] Kim, T., 2024. When is digital censorship permissible? A conversation norms account. Journal of Consumer Research doi:10.1093/jcr/ucae054.
- [106] King, G., Pan, J., Roberts, M.E., 2013. How censorship in china allows government criticism but silences collective expression. American political science Review 107, 326–343.
- [107] Kistenmacher, L., Talpur, A., Fischer, M., 2025. QUIC-aware load balancing: Attacks and mitigations, in: Proc. 55th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN '25), pp. 524–536. doi:10.1109/DSN64029.2025.00056.
- [108] Knockel, J., Crete-Nishihata, M., Ng, J.Q., Senft, A., Cran-

- dall, J.R., 2015. Every rose has its thorn: Censorship and surveillance on social video platforms in china, in: Proc. FOCI '15.
- [109] Knockel, J., Ruan, L., 2021. Measuring qqmail's automated email censorship in china, in: Proc. FOCI '21, ACM. p. 8–15. URL: <https://doi.org/10.1145/3473604.3474560>, doi:10.1145/3473604.3474560.
- [110] Knockel, J., Ruan, L., Crete-Nishihata, M., 2017. Measuring decentralization of Chinese keyword censorship via mobile games, in: Proc. FOCI '17. URL: <https://www.usenix.org/conference/foci17/workshop-program/presentation/knockel>.
- [111] Kubin, E., von Sikorski, C., Gray, K., 2024. Political censorship feels acceptable when ideas seem harmful and false. *Political Psychology* doi:<https://doi.org/10.1111/pops.13011>.
- [112] Kührer, M., Hupperich, T., Bushart, J., Rossow, C., Holz, T., 2015. Going wild: Large-scale classification of open DNS resolvers, in: Proc. 2015 Internet Measurement Conf. (IMC), ACM. pp. 355–368.
- [113] Lange, F., Niere, N., von Niessen, J., 2024. Censors ignore unencrypted HTTP/2 traffic. URL: <https://upb-syssec.github.io/blog/2024/http2/>.
- [114] Li, X., Li, D., Newman, J., 2013. Parental behavioral and psychological control and problematic Internet use among Chinese adolescents: The mediating role of self-control. *Cyberpsychology, Behavior, and Social Networking* 16, 442–447.
- [115] Liang, X., Liu, G., Jin, L., Hao, S., Wang, H., 2024. Pathfinder: Exploring path diversity for assessing internet censorship inconsistency. arXiv preprint arXiv:2407.04213 .
- [116] Lipphardt, F., Feldmann, A., Gosain, D., 2025a. 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. pp. 720–736.
- [117] Lipphardt, F., Tashiro, M., Fontugne, R., 2025b. 1-800-Censorship: Analyzing Internet censorship data using the Internet Yellow Pages, in: Proc. FOCI '25.
- [118] Lorimer, A.H., Tulloch, L., Bocovich, C., Goldberg, I., 2021. OUSTralopithecus: Overt user simulation for censorship circumvention, in: Proc. 20th Workshop on Workshop on Privacy in the Electronic Society (WPES'21), pp. 137–150.
- [119] Luo, Z., 2024. Negotiating censorial power and its legitimacy: a case study of the second face of state censorship. *Journal of Political Power* 17, 127–148. URL: <https://doi.org/10.1080/2158379X.2024.2370813>, doi:10.1080/2158379X.2024.2370813.
- [120] MacKinnon, R., 2011. China's "networked authoritarianism". *Journal of Democracy* 22, 32–46. doi:10.1353/jod.2011.0033.
- [121] Marder, B., Joinson, A., Shankar, A., Houghton, D., 2016. The extended 'chilling' effect of Facebook: The cold reality of ubiquitous social networking. *Comp. Human Behavior* 60, 582–592. doi:10.1016/j.chb.2016.02.097.
- [122] Margolis, D., Risher, M., Ramakrishnan, B., et al., 2018. SMTP MTA strict transport security (MTA-STS) - RFC 8461. <https://datatracker.ietf.org/doc/html/rfc8461>.
- [123] Marwick, A., Boyd, D., 2018. Understanding privacy at the margins. *International Journal of Communication* 12, 9.
- [124] Master, A., Garman, C., 2023. A worldwide view of nation-state Internet censorship, in: Proc. FOCI '23.
- [125] Maxmind Inc., 2025. IP geolocation and online fraud prevention website. <https://www.maxmind.com/>.
- [126] Mazurczyk, W., Wendzel, S., Cabaj, K., 2018. Towards deriving insights into data hiding methods using pattern-based approach, in: Proc. 13th International Conf. Availability, Reliability and Security (ARES 2018, ACM. pp. 10:1–10:10.
- [127] Mazurczyk, W., Wendzel, S., Chourib, M., Keller, J., 2019. Countering adaptive network covert communication with dynamic wardens. *Future Gen. Computer Systems (FGCS)* 94, 712–725.
- [128] Mazurczyk, W., Wendzel, S., Zander, S., Houmansadr, A., Szczypiorski, K., 2016. Information Hiding in Communication Networks: Fundamentals, Mechanisms, and Applications. Wiley-IEEE.
- [129] Meserve, S.A., Pemstein, D., 2018. Google politics: The political determinants of Internet censorship in democracies. *Political Science Research and Methods* 6, 245–263. doi:10.1017/psrm.2017.1.
- [130] Meserve, S.A., Pemstein, D., 2020. Terrorism and Internet censorship. *Journal of peace research* 57, 752–763.
- [131] Miao, W., Liu, J., Wu, S., 2024. Embedded symbiosis: an institutional approach to government-business relationships in the Chinese Internet industry, in: *The Geopolitics of Chinese Internets*. Routledge, pp. 113–130.
- [132] Mohajeri Moghaddam, H., Li, B., Derakhshani, M., Goldberg, I., 2012. SkypeMorph: Protocol obfuscation for Tor bridges, in: Proc. 2012 ACM Conf. Computer and Communications Security, ACM. pp. 97–108.
- [133] Mohd Aminuddin, M.A.I., Zaaba, Z.F., Samsudin, A., Zaki, F., Anuar, N.B., 2023. The rise of website fingerprinting on Tor: Analysis on techniques and assumptions. *Journal of Network and Computer Applications* 212, 103582. URL: <https://www.sciencedirect.com/science/article/pii/S1084804523000012>, doi:<https://doi.org/10.1016/j.jnca.2023.103582>.
- [134] Moriano, P., Hill, R., Camp, L.J., 2021. Using bursty announcements for detecting BGP routing anomalies. *Computer Networks* 188, 107835. URL: <https://www.sciencedirect.com/science/article/pii/S1389128621000207>, doi:10.1016/j.comnet.2021.107835.
- [135] Müller, P., Niere, N., Lange, F., Somorovsky, J., 2024. Turning attacks into advantages: Evading HTTP censorship with HTTP request smuggling. *Proc. Privacy Enhancing Technologies* .
- [136] Murdoch, S., Roberts, H., 2013. Introduction to Internet censorship and control. <https://murdoch.is/papers/ieeeic13control.pdf>.
- [137] Nabi, Z., 2013. The anatomy of web censorship in Pakistan, in: Proc. FOCI '13.
- [138] Nasr, M., Zolfaghari, H., Houmansadr, A., Ghafari, A., 2020. MassBrowser: Unblocking the censored web for the masses, by the masses., in: NDSS.
- [139] Networ, M.D., 2015. Same origin policy. URL: https://developer.mozilla.org/en-US/docs/Web/JavaScript/Same_origin_policy_for_JavaScript.
- [140] Niaki, A.A., Cho, S., Weinberg, Z., Hoang, N.P., Razaghpanah, A., Christin, N., Gill, P., 2020. ICLab: A global, longitudinal Internet censorship measurement platform, in: Proc. IEEE Symp. Security and Privacy (S&P '20), pp. 214–230. doi:10.1109/SP.2020.00014.
- [141] Niere, N., Lange, F., Merget, R., Somorovsky, J., 2025. Transport Layer Obscurity: Circumventing SNI Censorship on the TLS-Layer, in: 2025 IEEE Symposium on Security and Privacy (S&P), IEEE Computer Society. pp. 1288–1306. doi:10.1109/SP61157.2025.00151.
- [142] Nikiforakis, N., Younan, Y., Joosen, W., 2010. HProxy: Client-side detection of SSL stripping attacks, in: *Detection of Intrusions and Malware, and Vulnerability Assessment*, Springer Berlin Heidelberg. pp. 200–218. doi:10.1007/978-3-642-14215-4_12.
- [143] Nisbet, E.C., Kamenchuk, O., Dal, A., 2017. A psychological firewall? risk perceptions and public support for online censorship in russia. *Social Science Quarterly* 98, 958–975.
- [144] Noelle-Neumann, E., 1974. The spiral of silence a theory of public opinion. *Journal of Communication* 24, 43–51. doi:<https://doi.org/10.1111/j.1460-2466.1974.tb00367.x>.
- [145] Nourin, S., Tran, V., Jiang, X., Bock, K., Feamster, N., Hoang, N.P., Levin, D., 2023. Measuring and evading turkmenistan's internet censorship: A case study in large-scale measurements of a low-penetration country, in: Proc. ACM Web Conf.2023, pp. 1969–1979.
- [146] OONI, IODA, M-Lab, Cloudflare, Kentik, Censored Planet, ISOC, and Article19, 2022. Technical multi-stakeholder report on Internet shutdowns: The case of Iran amid au-


- tumn 2022 protests. <https://ooni.org/post/2022-iran-technical-multistakeholder-report/>.
- [147] OONI Project, 2024. Website. URL: <https://ooni.org/partners/ioda>.
- [148] OONI Project, 2025a. OONI measurement aggregation toolkit (MAT). URL: <https://explorer.ooni.org>.
- [149] OONI Project, 2025b. OONI probe tools. URL: <https://ooni.org/get-involved/>.
- [150] OONI Project, 2025c. Website. URL: <https://ooni.org>.
- [151] OpenNet Initiative, 2013. dataset website. URL: <https://opennet.net/research/data>.
- [152] OpenNet Initiative, 2014. project website. URL: <https://opennet.net/>.
- [153] Padmanabhan, R., Filastò, A., Xynou, M., Raman, R.S., Middleton, K., Zhang, M., Madory, D., Roberts, M., Dainotti, A., 2021. A multi-perspective view of Internet censorship in myanmar, in: Proc. FOCI '21, ACM. p. 27–36.
- [154] Panchenko, A., Niessen, L., Zinnen, A., Engel, T., 2011. Web-site fingerprinting in onion routing based anonymization networks, in: Proc. 10th Annual ACM Workshop on Privacy in the Electronic Society, ACM. pp. 103–114.
- [155] Paxson, V., 1998. On calibrating measurements of packet transit times, in: Proc. SIGMETRICS'98, ACM. pp. 11–21.
- [156] Paxson, V., 2004. Strategies for sound Internet measurement, in: Proc. IMC'04, ACM. p. 263–271. doi:10.1145/1028788.1028824.
- [157] Pearce, P., Jones, B., Li, F., Ensafi, R., Feamster, N., Weaver, N., Paxson, V., 2017. Global measurement of DNS manipulation, in: Proc. USENIX Security '17.
- [158] Penney, J.W., 2016. Chilling effects: Online surveillance and Wikipedia use. Berkeley Tech. LJ 31, 117.
- [159] Poell, T., 2014. Social media activism and state censorship, in: Social media, politics and the state. Routledge, pp. 189–206.
- [160] Polverini, B., Pottenger, W., 2011. Using clustering to detect chinese censorship, in: Proc. Seventh Annual Workshop on Cyber Security and Information Intelligence Research, ACM. doi:10.1145/2179298.2179331.
- [161] Psiphon Data Engine, 2024. Psiphon data engine. URL: <https://psix.ca/d/g6Mhku6Zz/faqs?orgId=2>.
- [162] Quan, L., Heidemann, J., Pradkin, Y., 2013. Trinocular: understanding internet reliability through adaptive probing, in: Proc. ACM SIGCOMM 2013 Conf. SIGCOMM, ACM. p. 255–266.
- [163] Radsch, C., 2023. Weaponizing privacy and copyright law for censorship. CIGI Papers No. URL: <http://dx.doi.org/10.2139/ssrn.4464300>. accessed: 2025-02-14.
- [164] Raman, R.S., Stoll, A., Dalek, J., Ramesh, R., Scott, W., Ensafi, R., . Measuring the Deployment of Network Censorship Filters at Global Scale, in: Proc. 2020 Network and Distributed System Security Symposium, Internet Society. URL: <https://www.ndss-symposium.org/wp-content/uploads/2020/02/23099.pdf>, doi:10.14722/ndss.2020.23099.
- [165] Raman, R.S., Wang, M., Dalek, J., Mayer, J., Ensafi, R., 2022. Network measurement methods for locating and examining censorship devices, in: Proc. 18th International Conf. Emerging Networking Experiments and Technologies, pp. 18–34.
- [166] Ravaioli, R., Urvoy-Keller, G., Barakat, C., 2015. Characterizing ICMP rate limitation on routers, in: 2015 IEEE International Conf. Communications (ICC), IEEE. pp. 6043–6049.
- [167] Report, G., 2021. The GFW has now been able to dynamically block any seemingly random traffic in real time. URL: https://x.com/gfw_report/status/1460796633571069955.
- [168] Rescorla, E., 2018. The Transport Layer Security (TLS) Protocol Version 1.3. RFC 8446. URL: <https://www.rfc-editor.org/info/rfc8446>, doi:10.17487/RFC8446.
- [169] Rescorla, E., Oku, K., Sullivan, N., Wood, C.A., 2024. TLS Encrypted Client Hello. Internet-Draft draft-ietf-tls-esni-22. Internet Engineering Task Force. URL: <https://datatracker.ietf.org/doc/draft-ietf-tls-esni/22/>. work in Progress.
- [170] RIPE, 2025. RIPE Atlas. URL: <https://atlas.ripe.net>.
- [171] Robertson, J., Diab, A., Marin, E., Nunes, E., Paliath, V., Shakarian, J., Shakarian, P., 2017. Darkweb cyber threat intelligence mining. Cambridge University Press.
- [172] Rosenberg, R.S., 2001. Controlling access to the internet: the role of filtering. Ethics and Information Technology 3, 35–54.
- [173] Ryng, J., Guicherd, G., Saman, J.A., Choudhury, P., Kellelt, A., 2022. Internet shutdowns. The RUSI Journal 167. doi:10.1080/03071847.2022.2156234.
- [174] Safaei Pour, M., Nader, C., Friday, K., Bou-Harb, E., . A comprehensive survey of recent internet measurement techniques for cyber security. Computers & Security 128, 103123. URL: <https://www.sciencedirect.com/science/article/pii/S0167404823000330>, doi:10.1016/j.cose.2023.103123.
- [175] Saputra, F.A., Nadhori, I.U., Barry, B.F., 2016. Detecting and blocking onion router traffic using deep packet inspection, in: 2016 International Electronics Symposium (IES), pp. 283–288. URL: <https://ieeexplore.ieee.org/abstract/document/7861018>, doi:10.1109/ELECSYM.2016.7861018.
- [176] Schlamp, J., Holz, R., Jacquemart, Q., Carle, G., Bierack, E.W., . HEAP: Reliable Assessment of BGP Hijacking Attacks 34, 1849–1861. URL: <https://ieeexplore.ieee.org/document/7460217?arnumber=7460217>, doi:10.1109/JSAC.2016.2558978.
- [177] Schulz, W.S., 2024. Warped Words: How Online Speech Misrepresents Opinion. Ph.D. thesis. Princeton Univ. URL: <http://arks.princeton.edu/ark:/88435/dsp01mc87pt570>.
- [178] Scott, W., Anderson, T., Kohno, T., Krishnamurthy, A., 2016. Satellite: Joint analysis of CDNs and Network-Level interference, in: 2016 USENIX Annual Technical Conf. (ATC 16), pp. 195–208.
- [179] Shapira, T., Shavitt, Y., 2022. AP2Vec: An Unsupervised Approach for BGP Hijacking Detection. IEEE Trans. Network and Service Management 19, 2255–2268. URL: <https://ieeexplore.ieee.org/document/9754706/>, doi:10.1109/TNSM.2022.3166450.
- [180] Sheffey, J., Houmansadr, A., . CensorLab: A Testbed for Censorship Experimentation. URL: <http://arxiv.org/abs/2412.16349>, doi:10.48550/arXiv.2412.16349, arXiv:2412.16349.
- [181] Simmons, G.J., 1983. The prisoners' problem and the subliminal channel, in: Proc. CRYPTO, pp. 51–67.
- [182] Subramanian, R., 2011. The growth of global Internet censorship and circumvention: A survey. Communications of the International Information Management Association (CIIMA) 11.
- [183] Sundara Raman, R., Shenoy, P., Kohls, K., Ensafi, R., 2020. Censored planet: An internet-wide, longitudinal censorship observatory, in: Proc. 2020 ACM SIGSAC Conf. Computer and Communications Security, ACM. p. 49–66. URL: <https://doi.org/10.1145/3372297.3417883>, doi:10.1145/3372297.3417883.
- [184] Syverson, P., Traudt, M., 2018. HSTS supports targeted surveillance, in: Proc. FOCI '18.
- [185] Tang, J., Alvarez, L., Brar, A., Hoang, N.P., Christin, N., 2024. Automatic generation of web censorship probe lists, in: Proc. Privacy Enhancing Technologies.
- [186] The Tor Project, 2025. Tor metrics. URL: <https://metrics.torproject.org>.
- [187] Tippe, P., Tippe, A., 2024. Onion services in the wild: A study of deanonymization attacks. Proc. Privacy Enhancing Technologies 2024 (PoPETs) 2024, 291–310. doi:10.56553/popets-2024-0117.
- [188] Tsai, E., Kumar, D., Raman, R.S., Li, G., Eiger, Y., Ensafi, R., 2023. CERTainty: Detecting DNS manipulation at scale using TLS certificates, in: Proc. Privacy Enhancing Technol.
- [189] Tsai, E., Raman, R.S., Prakash, A., Ensaf, R., 2024. Modeling and detecting internet censorship events, in: NDSS'24. doi:10.14722/ndss.2024.23409.
- [190] Tschantz, M.C., Afroz, S., anonymous, Paxson, V., 2016. SoK: towards grounding censorship circumvention in empiricism, in: Proc. IEEE Symp. Security and Privacy (S&P '16), pp. 914–

933. doi:10.1109/SP.2016.59.
- [191] VanderSloot, B., McDonald, A., Scott, W., Halderman, J.A., Ensafi, R., 2018. Quack: Scalable remote measurement of application-layer censorship, in: Proc. USENIX Security Symp. '18.
- [192] Ververis, V., Marguel, S., Fabian, B., 2020. Cross-country comparison of Internet censorship: A literature review. *Policy & Internet* 12, 450–473.
- [193] Vines, P., McKay, S., Jenter, J., Krishnaswamy, S., 2024. Communication breakdown: Modularizing application tunneling for signaling around censorship. *Proc. Privacy Enhancing Technologies*.
- [194] Wails, R., Sullivan, G.A., Sherr, M., Jansen, R., . On Precisely Detecting Censorship Circumvention in Real-World Networks, in: Proceedings 2024 Network and Distributed System Security Symposium, Internet Society. URL: <https://www.ndss-symposium.org/wp-content/uploads/2024-394-paper.pdf>, doi:10.14722/ndss.2024.23394.
- [195] Wampler, J., Tanveer, H.B., Nithyanand, R., Wustrow, E., 2025. ProtoScan: Measuring censorship in IPv6. arXiv preprint arXiv:2508.07194.
- [196] Wang, D., Mark, G., 2015. Internet censorship in China: Examining user awareness and attitudes. *ACM Trans. Computer-Human Interaction (TOCHI)* 22, 1–22.
- [197] Wang, Z., Cao, Y., Qian, Z., Song, C., Krishnamurthy, S.V., 2017. Your state is not mine: a closer look at evading stateful Internet censorship, in: Proc. IMC '17, ACM. p. 114–127. doi:10.1145/3131365.3131374.
- [198] Warf, B., 2011. Geographies of global Internet censorship. *GeoJournal* 76, 1–23. doi:10.1007/s10708-010-9393-3.
- [199] Weinberg, Z., Wang, J., Yegneswaran, V., Briesemeister, L., Cheung, S., Wang, F., Boneh, D., 2012. StegoTorus: a camouflage proxy for the tor anonymity system, in: Proc. 2012 ACM Conf. Computer and Communications Security (CCS'12), ACM. pp. 109–120.
- [200] Wendzel, S., Caviglione, L., Mazurczyk, W., 2023. Avoiding research tribal wars using taxonomies. *Computer* 56. doi:10.1109/MC.2022.3218175.
- [201] Wendzel, S., Caviglione, L., Mazurczyk, W., Mileva, A., Dittmann, J., Krätzer, C., Lamshöft, K., Vielhauer, C., Hartmann, L., Keller, J., Neubert, T., Zillien, S., 2025. A generic taxonomy for steganography methods. *ACM Computing Surveys (CSUR)* 57. doi:10.1145/3729165.
- [202] Wendzel, S., Zander, S., Fechner, B., Herdin, C., 2015. Pattern-based survey and categorization of network covert channel techniques. *ACM Computing Surveys (CSUR)* 47. doi:10.1145/2684195.
- [203] Winter, P., 2014. Measuring and Circumventing Internet Censorship. Ph.D. thesis, Karlstad University. URL: <https://urn.kb.se/resolve?urn=urn:nbn:se:kau:diva-34475>.
- [204] Winter, P., Pulls, T., Fuss, J., 2013. ScrambleSuit: A polymorphic network protocol to circumvent censorship, in: Proc. 12th ACM workshop on Privacy in the electronic society, pp. 213–224.
- [205] Wrana, M., Barradas, D., Asokan, N., 2025. SoK: The spectre of surveillance and censorship in future internet architectures. *Proc. Privacy Enhancing Technologies*.
- [206] Wu, M., Sippe, J., Sivakumar, D., Burg, J., Anderson, P., Wang, X., Bock, K., Houmansadr, A., Levin, D., Wustrow, E., 2023. How the great firewall of china detects and blocks fully encrypted traffic, in: Proc. USENIX Security Symp. '23.
- [207] Wu, M., Zohaib, A., Durumeric, Z., Houmansadr, A., Wustrow, E., 2025. A wall behind a wall: Emerging regional censorship in china, in: 2025 IEEE Symposium on Security and Privacy (S&P), IEEE Computer Society. pp. 1307–1324. URL: <https://doi.ieeecomputersociety.org/10.1109/SP61157.2025.00152>, doi:10.1109/SP61157.2025.00152.
- [208] Wübbeling, M., Sykosch, A., M.Meier, 2017. Quantum suite - a look inside the NSA toolbox, in: Tagungsband zum 15. Deutschen IT-Sicherheitskongress, pp. 239–254.
- [209] Xie, Y., Gou, G., Xiong, G., Li, Z., Xia, W., 2024. Dom-
Eye: Detecting network covert channel of domain fronting with throughput fluctuation. *Computers & Security* 144. doi:10.1016/j.cose.2024.103976.
- [210] Xue, D., Ramesh, R., Jain, A., Kallitsis, M., Halderman, J.A., Crandall, J.R., Ensafi, R., 2024. OpenVPN is open to VPN fingerprinting 68. URL: <https://dl.acm.org/doi/10.1145/3618117>, doi:10.1145/3618117.
- [211] Xue, D., Ramesh, R., S, V.S., Evdokimov, L., Viktorov, A., Jain, A., Wustrow, E., Basso, S., Ensafi, R., 2021. Throttling Twitter: An emerging censorship technique in Russia, in: Proc. 21st ACM Internet Measurement Conference, ACM. pp. 435–443. URL: <https://dl.acm.org/doi/10.1145/3487552.3487858>, doi:10.1145/3487552.3487858.
- [212] Zander, S., Armitage, G., Branch, P., 2007. A survey of covert channels and countermeasures in computer network protocols. *IEEE Communications Surveys & Tutorials* 9, 44–57.
- [213] Zittrain, J.L., Faris, R., Noman, H., Clark, J., Tilton, C., Morrison-Westphal, R., 2017. The shifting landscape of global Internet censorship.
- [214] Zook, M.A., Graham, M., 2007. The creative reconstruction of the Internet: Google and the privatization of cyberspace and DigiPlace. *Geoforum* 38, 1322–1343. doi:<https://doi.org/10.1016/j.geoforum.2007.05.004>.
- [215] Zuboff, S., 2019. The age of surveillance capitalism. Profile Books.
- [216] Zuchora-Walske, C., 2010. Internet Censorship: Protecting Citizens or Trampling Freedom? Lerner Publishing Group.

Author Biographies



Steffen Wendzel is a professor at Ulm University, Germany, where he is the chair of the Institute of Information Resource Management (IRM) and the director of the university computing center and university library (kiz). Previously, he was a professor at HS Worms, Germany, a lecturer (*Privatdozent*) at the Dep. Mathematics & Computer Science, FernUniversität in Hagen, Germany, and a PostDoc at Fraunhofer FKIE, Bonn, Germany. He received his PhD (2013) and Habilitation (2020) from the University of Hagen, Germany. Steffen (co-)authored more than 190 publications, of which many appeared in major journals and conferences (e.g., Elsevier COSE, Elsevier FGCS, ACM CSUR, IEEE TDSC, IEEE S&P Mag., AsiaCCS, IEEE LCN, Comm. ACM, ARES etc.). Website: <https://www.wendzel.de> / <https://www.uni-ulm.de/en/in/omi/>.



Simon Volpert is a researcher at the Institute of Information Resource Management (IRM) at the University of Ulm and PhD student at the UzK, Germany. He received his M.Sc. at Ulm University in 2023. His research interests revolve around the challenges of operations at scale, as well as systems benchmarking and data processing. Website: <https://www.uni-ulm.de/en/in/omi/>.



Sebastian Zillien is a researcher at the Institute of Information Resource Management (IRM) at the University of Ulm and PhD student at the RPTU, Germany. Earlier, he worked for the projects WoDiCoF+ and ATTRIBUT at

the University of Applied Sciences in Worms, Germany. He received his M.Sc. in Mobile Computing in 2020 from Hochschule Worms. His research focus includes network and protocol-level security, anomaly detection, covert channels and reliability. Website: <https://www.uni-ulm.de/en/in/omi/>.



Julia Lenz is a PhD student at the Institute of Information Resource Management (IRM) at the University of Ulm. She received her B.Sc. from the University of Essex, UK, in Psychology and her M.Sc. from the City University of

London, UK, in 2020 with the specialization in Organizational Psychology. Thereafter, she worked on several projects as a scientific researcher at the University of Applied Science in Worms, Germany. With her background in Psychology, her research interests include human factors in cybersecurity and usable security. Website: <https://www.uni-ulm.de/en/in/omi/>.



Philip Rünz has been a Policy Advisor for Security and Defense in the German Bundestag since 2023. He completed a dual study program in Mechatronics at the Baden-Württemberg Cooperative State University (DHBW) in

cooperation with the German Armed Forces (B.Eng., 2020) and later specialized in cybersecurity with a degree in Practical Computer Science from the FernUniversität Hagen (M.Sc., 2023). From 2021 to 2023, he worked on a cyber project for the Bundeswehr. Since 2021, he has also been lecturing on signal and control technology at the DHBW Mannheim.



Luca Caviglione is a Senior Research Scientist at the Institute for Applied Mathematics and Information Technologies of the National Research Council of Italy. He holds a Ph.D. in Electronic and Computer Engineering from

the University of Genoa, Italy. His research interests include optimization of large-scale computing and network security. He is an author or co-author of more than 200 academic publications, and several patents in the field of p2p and energy-aware computing. He is an associate editor of the IEEE Transactions on Information Forensics & Security and the head of the IMATI Research Unit of the National Inter-University Consortium for Telecommunications. Website: <https://lucacav.github.io>.