The WorldWideBlock Toolkit: Tracking Public Blocklists and Active Censorship Across 190 Countries

Jan Riedler, Vasilis Ververis, and Vaibhav Bajpai
Hasso Plattner Institute and University of Potsdam, Germany
{jan.riedler, vasilis.ververis, vaibhav.bajpai}@hpi.de

***Abstract*—This paper presents the first global analysis of blocked domains at the country and ISP level. It examines the transparency of governments and ISPs regarding the domains they block and compares this to the actual censorship observed in practice. To assess this transparency, 40 officially published blocklists by governments and ISPs are automatically crawled regularly, enabling both cross-country comparisons and longitudinal evaluation over time. These blocklists are also made publicly available through the developed platform *WorldWideBlock*. To identify the domains that are actually blocked, over 20 million Open Observatory of Network Interference (OONI) anomalies are analyzed and validated through more than 500 million conducted remote measurements, enabling the extraction of anomalies caused specifically by ISP- or country-level blocking. Through this comparison, domain blocklists were generated for 190 countries, enabling cross-country comparisons of censorship strategies and revealing inconsistencies in transparency and enforcement practices. Additionally, our contributions include the development of two measurement tools, rVPmt and rDNSmt, which enhance our ability to analyze and understand Internet censorship dynamics. Finally, the extracted blocklists were applied to RIPE Atlas measurements to identify the share of tested domains that could potentially harm probe hosts due to censorship-related redirects or interception.**

***Index Terms*—Censorship, Blocklist, Blacklist, Government, ISP, OONI, Censored Planet, Block, Web**

I. Introduction

The Internet has transformed communication, offering users unprecedented freedom of expression and access to information, as envisioned by Tim Berners-Lee [1]. However, in recent decades, governments have increasingly restricted this freedom by blocking access to websites based on their interpretations of rights. For instance, since February 2025, the Italian communications authority AGCOM has operated the "Piracy Shield" system, an automated system enabling copyright holders to request the rapid blocking of websites and IP addresses hosting illegal content within 30 minutes [2]. Reports from organizations like Freedom House indicate a steady global rise in state-imposed restrictions, limiting online communication and information access [3]. Internet Service Providers (ISPs) play a crucial role in implementing these restrictions, often through curated blocklists, meaning lists of websites or services that are restricted. While some governments disclose parts of these blocklists, many do not [4]. Therefore, not only the access to information is restricted, but also the scope and nature of the censorship itself. Without public knowledge of what content is being restricted, it becomes difficult to assess whether such measures are legally and ethically justified or whether they function primarily as instruments of political control and suppression of dissent. Therefore, understanding the scope and mechanisms of blocklists is essential for advancing our knowledge of digital freedom.

Previous studies on domain blocking have primarily focused on specific topics, regions, or particular censorship techniques (Section II-A). To this end, they often relied on the OONI dataset or conducted their own remote measurements, such as those performed by Censored Planet (see Section II for a background on these platforms). Both approaches face the challenge that detected blocking events cannot always be clearly attributed to the responsible network operator and therefore require more detailed analyses and comparisons (Section II). This approach improves the clear identification of active network-operator blocking by comparing local OONI measurements with remote measurements and reviewing published information. To date, no comprehensive global assessment of authoritative, publicly released blocklists exists. This gap motivates the following research question:

RQ 1: Which blocklists defined by network regulators are publicly accessible? This inquiry examines the availability of information on national and ISP blocklists, focusing on publicly accessible lists from governments and regulatory bodies. Analyzing these publications reveals how informative countries are to their citizens and facilitates comparisons between nations.

RQ 2: To what extent are domains blocked in practice? This exploration aims to identify domains blocked by ISPs or governments, providing a global view of domain censorship and the overall state of Internet freedom.

RQ 3: What techniques are currently implemented by network operators to enforce blocklists? This investigation examines the blocking techniques employed by network operators and how these methods vary across different categories of providers, based on the findings from RQ 2.

This paper aims to provide an up-to-date overview of blocklists in terms of their scope, implementation, and public availability. In addition to the systematic search for authoritative blocklists, it includes a large-scale analysis of OONI data combined with validation through remote measurements

ISBN 978-3-903176-82-9 © 2026 IFIP

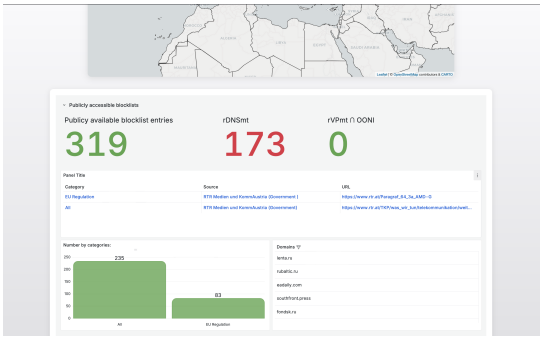

Fig. 1. Screenshot of the "worldwideblock" website. The public authoritative blocklists for Austria are shown, including the number of entries with the option to view the entries directly. Additionally, the number of confirmed OONI anomalies for rDNSmt and rVPmt measurements is displayed.

in order to more accurately identify blockades enforced by network regulators (Section III).

The main contributions are:

- **Development of WorldWideBlock:** An automated platform that systematically crawls 40 officially published government and ISP blocklists. This open-source tool enables longitudinal monitoring and public visualisation of both official restrictions and confirmed active blocking, addressing a critical gap in global transparency.
- **New Remote Measurement Methods:** We introduce two new tools, `rVPmt` and `rDNSmt`, designed to improve the accuracy of censorship attribution. By performing more than 500 million remote measurements to validate over 20 million OONI anomalies, these techniques distinguish between ISP-level censorship and false positives caused by local firewalls or server-side errors.
- **Global Mapping of Censorship Enforcement:** We present a comprehensive blocklist for 190 countries, providing a global characterisation of censorship strategies. Our findings reveal that while DNS-based techniques (such as `NXDOMAIN` responses or unroutable IPs) are nearly universal, some nations also employ `Server Name Indication (SNI)` and `HTTP` host header interception to enforce content restrictions.
- **Practical Safety Applications for RIPE Atlas Hosts:** We apply our global blocklists to RIPE Atlas to estimate the proportion of measurements that may target sensitive or censored domains. We demonstrate how authoritative blocklists can serve as a safeguard for probe hosts, protecting them from potential risks associated with censorship-related redirects or interception.

See Section VI for a more detailed discussion of the research contributions. All collected data, analysis scripts, and the WorldWideBlock platform will be released under a free software license, ensuring transparency and reproducibility upon acceptance of this paper.

II. Background and Related Work

In this paper, “Blocked domains” refer to those intentionally restricted by ISPs or governments. See RFC 7754 for further technical details on domain blocking and filtering [5]. The two largest platforms, providing open-access global data on domain accessibility, are OONI (Open Observatory of Network Interference) and Censored Planet [6]. OONI detects and documents network interference through a volunteer network conducting tests via the Web Connectivity experiment, which checks domain accessibility through DNS lookups, TCP connections, and HTTP(S) requests [4], [7]. Discrepancies between results from different networks mark a URL as anomalous. However, OONI cannot identify the source of the block (ISP, local settings, or external restrictions), leading to potential misclassification of domains as blocked by specific ISPs. As an example, an anomaly was observed for chatgpt.com from AS3320 [8], which manual verification revealed to be a false positive. Effectively filtering out these misclassifications is both difficult and central to this approach. An alternative procedure to identifying which domains are blocked by ISPs are remote measurements, where tests are conducted from a controlled measurement point to external vantage points (VP). Censored Planet implements this technique through its tools Satellite and Hyperquack [9], which operate on a global scale. Satellite evaluates discrepancies between open and control resolvers, whereas Hyperquack detects blocking mechanisms in HTTP(S) traffic based on request headers. Both tools, are closed-source and not publicly available. OONI and Censored Planet rely on Citizen Lab’s global and regional test lists as the foundation for selecting the domains to be analyzed. Additionally, OONI allows users to test their own domains and Censored Planet is adding some domains from Alexa’s Top List, which lists popular domains. While Censored Planet tests 2000 domains weekly from these lists, OONI has measurements on more than 27000 different domains [10], [11]. Consequently, the data from Censored Planet cannot be used to cross-validate all the OONI data, although remote measurements would, in principle, be suitable for filtering out ISP/government-induced blocks (see Section III).

A. Related Work

Recent studies highlight the limitations of existing datasets in identifying all instances of ISP-imposed blocks. For example, Vyas et al. [12] conducted a global study on blocked COVID-19 websites using Censored Planet’s tools across 180 countries, emphasizing the need for comprehensive data. Similarly, Ververis et al. [4] analyzed network interference in 27 EU countries using OONI data, filtering out local blocks and validating instances of censorship. Their findings revealed that authorities issue various blocklists, often without clear justification for certain restrictions. While there have been analyses of OONI data, they tend to focus on specific topics, such as censorship of LGBTIQ content in Tanzania and Telegram blocking in Brazil. This paper aims to update previous findings, enhance the validation process of OONI data, and extend the analysis globally. While some studies have examined particular censorship techniques [13]–[15], a

comprehensive global evaluation comparing these methods across different provider categories is still lacking.

III. Methodology

To effectively address the research questions, the following approach was implemented: A web service, *WorldWideBlock*, was developed to automate the process of assessing and visualizing Internet censorship data. This service provides up-to-date insights into public blocklists, their contents and statistics reflecting the current state of censorship at the national level. Furthermore, all source code, data, and data analysis scripts related to *WorldWideBlock* will be released and made publicly available under a free software license upon acceptance of this paper. This commitment ensures that researchers and developers can access, modify, and contribute to the project, fostering collaboration and innovation in the field of Internet censorship analysis.

Identification of Public Authoritative Blocklists: Systematic searches were conducted using Google and platforms like TorrentFreak, NTC.Party, Freedom House, Citizen Lab, and Access Now. The search query structure is: `[country] internet (blocklist OR blacklist OR block OR list OR forbidden OR censorship OR filter OR illegal OR gov OR authority) [category]`. Categories include Gambling, Tobacco, Media, EU Regulation, Copyright, Courts, and Pornography, based on Ververis et al. [4]. Keywords for official websites were translated into local languages using DeepL [16]. All identified blocklists are downloaded daily and stored in a database for longitudinal analysis, with each entry including the domain/IP, start date, last observed date, category, source, and URL. Python was used for data extraction, and SQLite served as the database system for its simplicity. To bypass bot detection, requests were made using Selenium for browser automation. Blocklists are available in various formats, including TXT, JSON, CSV, Excel, and PDF. A general Python function processes these formats using libraries like pandas for Excel and PyPDF2 for PDFs. Regular expressions identify domains and dates for accurate data entry, while manual specifications are used for exclusions. For HTML tables, Selenium extracts relevant data directly. This methodology ensures efficient data collection and analysis while maintaining ethical standards.

Censorship Block Detection: Based on OONI data, it is generally not possible to determine who is responsible for anomalies unless a blockpage is returned. In the case of DNS-based anomalies, it cannot be determined whether the incorrect response originates from the resolver itself, a censorship device intercepting the DNS request or from natural network errors. Even non-DNS anomalies may result from various causes, including local firewalls (e.g., parental control or network policies), server-side restrictions (e.g., geoblocking), network errors, or actual blocks by a network operator. To identify OONI measurements that likely result from ISP- or country-level censorship devices, remote measurement techniques were implemented for cross-verification. For non-DNS blocks, such as those potentially caused by HTTP host header or SNI interception, the goal of the technique is to send a message containing the same test domain as in the OONI anomaly past the potential censorship devices and observe whether any interference occurs. If no interference is detected, it can be assumed that the domain was not actually blocked by network operators. To send the message past the potential censorship devices, so-called VP were used as IP targets, while the header information specified the test domains as the intended destination. For this, the VP must be located in the same Autonomous System (AS) as the machine that triggered the OONI anomaly. If both measurements, the OONI test and the remote measurement, indicate blocking, the likelihood that the blocking was caused by a network operator increases significantly. However, individual false positives may still occur in scenarios where the local network policies of the OONI probe and the security systems of the VP independently block the same domains (see Figure 2). For OONI anomalies that are based on DNS blocks, the used resolver is retested or used as VP if it is unreachable to determine whether a censorship device intercepts connections for specific test domains.

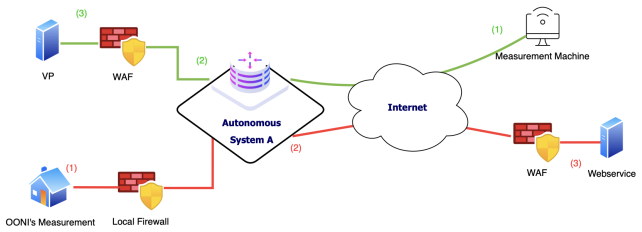

Fig. 2. Cross-validating OONI measurements with our own remote measurements to identify blocks caused by network operators (for non-DNS blocks). The figure illustrates that an existing OONI anomaly may result from multiple points along the web connectivity path: a local firewall, the connected AS, or the webservice protected by a Web Application Firewall (WAF). Each of these points can be responsible for the observed anomaly. To verify whether the reported AS is blocking a domain, our measurement machine sends a corresponding signal to a VP within the same AS and checks if the domain is blocked. If both measurements indicate an anomaly, it is likely that the censorship device within the AS is responsible.

Remote Measurements Procedure: To perform remote measurements rVPmt (remote vantage point measurement technique) and rDNSmt (remote DNS measurement technique) were created. Upon request, Censored Planet provided the VPs used for Hyperquack, which uses regular web servers as VPs, not custom probes. These servers, operated by ISPs or government entities, were categorized into CSV files based on supported protocols (Echo, HTTP, HTTPS). The analysis utilized 20586 Echo-, 11312 HTTP-, and 14732 HTTPS-servers.

rVPmt: To compare remote measurements with OONI data, all domains tested by OONI within a specific AS were also test using rVPmt with a VP in the same AS. For broader coverage, rVPmt utilized a list of distinct domains aggregated across all ASes, selecting a single reachable VP from each AS to test the

entire domain list. To initiate rVPmt, domains and VPs must be defined, including their IP addresses and types (ECHO, HTTP, HTTPS). The control domain and maximum timeout (default: 10 seconds) can also be adjusted. The procedure checks VP reachability and response times, retaining the fastest VP for each protocol type. For each selected VP, all domains are tested for blocking by sending a request to the VP using the control domain as the spoofed target. If the test value differs from the control value, it indicates potential censorship. Errors are categorized as censorship-related or natural, with retries until three errors of either type occur. Valid responses are compared, and results are stored in a JSONL file, categorized by protocol.

Requesting Comparable Values via HTTP: For VPs supporting HTTP, an HTTP GET request is sent to the IP address with the target domain in the `Host` header. The HTML body length of the response is extracted as the control value, after removing the target domain from the body if it is present. A differing body length suggests a blockpage, and the HTML title is extracted for verification.

Requesting Comparable Values via HTTPS: For VPs supporting HTTPS, the TLS handshake is performed with the target domain set in the SNI. The resulting certificate chain is validated against trusted root authorities. If valid, the result is marked as `"valid"`; otherwise, the raw leaf certificate is stored. If no certificate is retrieved, the result is marked as `"no certs"`.

Requesting Comparable Values via ECHO: Since the ECHO protocol simply mirrors the received content, the comparison value is the original request itself. To perform the measurement, an HTTP GET request is sent to the server with the target domain specified in the host header. If the response matches the request exactly, no censorship is assumed. However, if the response differs from the request, it can be assumed that the response came from a censorship device. A similar approach is taken using a TLS Client Hello, where the target domain is inserted into the SNI field. The response is then checked to determine whether the same Client Hello is returned, allowing detection of manipulation at the TLS layer.

rDNSmt processes OONI anomalies where DNS was indicated as the blocking reason. The first request is querying a control domain using the test resolver, the resolver used in the OONI test, to verify its reachability and functionality. On error, the domain being tested is still sent via UDP. A response differing from the control indicates that a censorship device intercepted the request and provided an answer. If the resolver was reachable, the test domain is queried using a control resolver (e.g., Cloudflare DNS: 1.1.1.1) to ensure it can be resolved. Afterward, the test domain is queried up to three times via the test resolver. Valid responses are classified as either an IP address, an Extended DNS Error (EDE) indicating blocking (15-18), or the standard DNS error code NXDOMAIN. IP addresses are further analyzed by comparing them with the IP addresses returned by the control resolver. If a valid, blocking-related error is returned, the result is marked as censored. In the cases of invalid error messages, the control domain is queried again to verify whether the resolver is still responsive. The IP address comparison is performed as follows:

If the test resolver returns non-routable IPs (e.g. 127.0.0.0/8, 10.0.0.0/8, 0.0.0.0) while the control resolver provides a public one, the result is marked as censored. If public IPs are returned, TLS certificates are validated. If validation fails, the certificate returned by the control resolver is compared to determine potential censorship, indicating the presence of a blockpage. If no TLS certificate is retrieved but a TCP connection to port 80 is successfully established, the result is also marked as censored, again suggesting a blockpage. For each DNS request, rDNSmt first attempts to resolve to an IPv4 address. If the response is `NOERROR` with `NODATA`, an additional query retrieves the IPv6 address. Each request and result is cached to avoid redundant lookups. The output is written to a JSONL file, with each line corresponding to a specific domain-resolver pair. To address unreachable DNS resolvers, OONI's recorded DNS response data is used as a fallback, allowing consistent evaluation.

Analyzing OONI Data: Since OONI's Explorer and API do not support large-scale analyses, a local mirror of the data must first be created for custom analysis. ClickHouse provides native support for reading data directly from S3 buckets, enabling efficient storage of the required information from OONI's public S3 bucket `ooni-data-eu-fra`. The processed data contains nearly all the required information, except for the DNS resolver used. To obtain the resolver for each measurement, the raw data must be extracted and matched using the report ID. A CRON job triggers a shell script monthly to download the corresponding processed and raw data.

Data Validation: Both procedures, rVPmt and rDNSmt, conclude under certain conditions that a blockpage is presented. As both run the risk of overestimating, the detected blockpages must be validated. For this purpose, pattern sets from Censored Planet are used, which identify characteristic HTML signatures of known blockpages. First, each page is checked against a set of false-positive patterns. If none of these match, the known blockpage patterns are applied [17]. If neither a false-positive nor a blockpage pattern is detected, the result is marked as *No pattern matched*. Since the rVPmt measurements can also result in blockpages from the VP's security system, a distinction must be made between blockpages from companies and those from ISPs/governments. Since the Censored Planet patterns also have this subdivision in their patterns, they can also be used for this purpose. These patterns distinguish between different categories, including commercial products, national firewalls, ISPs filters, corporate networks, unknown sources, and general signatures.

While both rDNSmt and OONI DNS measurements are capable of identifying which DNS resolvers block specific domains, accurately interpreting ISP-level blocking requires

distinguishing between resolvers operated by ISPs (serving residential customers) and those that are not. However, no existing method or publicly available database was found that reliably maps resolver IP addresses to ISPs or third-party services. Developing a validated methodology for this purpose goes beyond the scope and timeframe of this paper. To still exclude some resolvers that offer blocking as a service (e.g., those that filter advertising, spam, or phishing domains) or resolvers that are clearly operated by companies and apply filtering based on corporate policies, the following approach is implemented: For each resolver IP address, a PTR DNS query (reverse DNS lookup) is performed, which attempts to resolve an IP address back to a hostname. If the resulting hostname contained any of a set of predefined substrings associated with non-residential ISP resolvers, the corresponding results were excluded from the final blocklist.

Blocklists from OONI Measurements: The final blocklist aims to compile entries likely representing genuine censorship by ISPs or governments. Three distinct sublists are defined: 1. The first list consists of measurements where rVPmt identified an ISP- or nation-level blockpage. 2. The second list merges OONI anomalies (where reason for a block is not DNS) and rVPmt censorship detections, including a domain only if marked as censored by both within the same ASN. 3. The third list includes domains flagged through the rDNSmt. After combining these lists, the results can be sorted by country and AS, revealing the corresponding blockpages for each region.

Applying Global Blocklists to RIPE Atlas Measurements: RIPE Atlas is not intended for censorship measurements and explicitly discourages targeting domains that could potentially pose a risk to probe hosts [18]. For this analysis, publicly available measurement data from September 2024 to February 2025 was collected and structured into a dataset containing the tested domains and their associated measurement types. These domains were then compared against global blocklists to identify targets that are known to be blocked or considered sensitive in certain countries. The goal of this analysis is to estimate the proportion of RIPE Atlas measurements that may inadvertently include censored or high-risk domains, thereby demonstrating how blocklists could support future safeguards within the RIPE Atlas platform.

IV. Results

This section presents the findings for each research question, following the procedure described in Section III. The published blocklists already collected indicate significant differences in the enforcement of these regulations, both in terms of the blocking mechanisms used and the number of domains affected. Since the analysis by Ververis et al. [4], the regulations have significantly intensified.

A. Publicly Available Blocklists

Figure 3 presents 40 identified official blocklists and their entry counts, accessible via *WorldWideBlock*. Sources are linked in the descriptions, with additional explanations for some: Kazakhstan's government offers an API to check

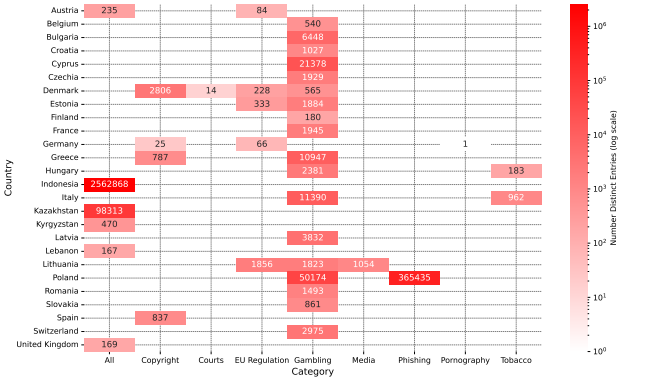

Fig. 3. The 40 authoritative, published blocklists found, grouped by country and category, and visualized by number of entries. A significant portion of the identified blocklists belong to the "Gambling" category, indicating that authorities demonstrate transparency primarily in this area.

blocked domains. Querying with a single letter returns all blocked domains containing that letter or an error. Using a period followed by a letter (e.g., .a, .b) successfully extracts the entire blocklist, including the block start date and IP address if applicable. The "Clearingstelle Urheberrecht im Internet" in Germany coordinates service blocking by recommending websites for copyright violations. These recommendations, including start dates, are published on their site, listing only one domain per service with a spoofed top-level domain. Indonesia's Ministry of Communications provides a tool to check if a domain is on the Trust+ blocklist, which includes inappropriate sites. A downloadable version of the blocklist obfuscates certain characters, while a publicly accessible DNS zone file contains the full list but is restricted to Indonesian IP addresses. Roskomnadzor allows users to check if a domain is blocked in Russia but does not publish the complete blocklist. Each query requires CAPTCHA verification, complicating full list extraction. Comparing blocklists from Ververis et al. [4] in 2020 to current data shows a consistent increase in blocked domains. Poland's gambling and phishing blocklists grew from approximately 14500 to over 50000 and 365000, respectively. Cyprus expanded its gambling blocklist from 13789 to over 21000, while Greece increased from around 2800 to nearly 10000. Denmark's count rose from fewer than 800 to over 3000 across various categories. Austria's blocked domains more than tripled from 85. Other countries like Italy, Romania, and Bulgaria also show upward trends, reflecting intensified state regulation of online content.

Several countries publish blocklists based on EU regulations, indicating that implementation is at the discretion of individual member states. The EU issues recommendations for blocking certain services, but enforcement varies significantly, both in scope and content (see Fig. 4). Lithuania and Estonia include the largest number of domains, even when only base domains are considered. Lithuania is the only country that also lists IP addresses in its blocklist. Austria and Denmark provide nearly identical lists, while Germany has the smallest number of entries. Merging these national lists could create

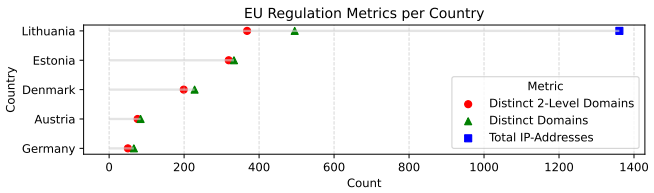

Fig. 4. Distribution of entries in publicly available blocklists based on EU regulations by country. Data are categorized by distinct second-level domains, total distinct domains, and IP addresses. Lithuania is the only country that includes IP addresses and contains the most entries, while Germany records the fewest.

a more comprehensive EU-level blocklist, enhancing content restriction effectiveness.

B. OONI

This evaluation covers the period from September 2024 to February 2025, during which 20496415 measurements were conducted across 180 countries using OONI's Web Connectivity test, resulting in numerous anomalies. Interestingly, no clear correlation with expected censorship patterns emerges.

For instance, France and Brazil are among the top three countries in terms of detected anomalies and also countries such as the Netherlands, Germany, and the United States also rank prominently, despite being classified as "free" according to Freedom House's Freedom on the Net report, with scores of 76/100 (France), 64/100 (Brazil), 94/100 (Netherlands), 76/100 (Germany), and 76/100 (United States), respectively [19]. In contrast, Kazakhstan, which is classified as "not free" with a score of 34/100, ranks only 25th in terms of detected anomalies. Nevertheless, in the top countries with the highest number of anomalies are also Russia, China, Venezuela, and Iran, which are all classified as not free in the Freedom on the Net report. Among them, Venezuela received the highest score, with 13 out of 100 points. This discrepancy may reflect the influence of a larger local measurement community in countries with many anomalies, leading to better coverage and more reliable analysis. In total, 136016 distinct domains were tested using OONI's Web Connectivity test, which can also be assessed with rVPmt. Among all OONI measurements, 7538056 anomalies were identified as DNS-based blocking, making them suitable for validation with rDNSmt.

C. rDNSmt

rDNSmt collected 7090259 measurements, including DNS results from OONI where resolvers were unreachable, covering over 94 % of OONI's DNS blocking indications. Most measurements involved unreachable resolvers, with 31105 out of 32217 tested servers inaccessible, leaving only 1112 reachable resolvers. This indicates that remote measurements are insufficient for evaluating global DNS blockage by resolvers.

Evaluation Excluding OONI's DNS Responses: Only distinct domain–resolver pairs with reachable resolvers were considered. 36127 pairs were classified as not censored, while 5381 were marked as censored. Most domains were classified

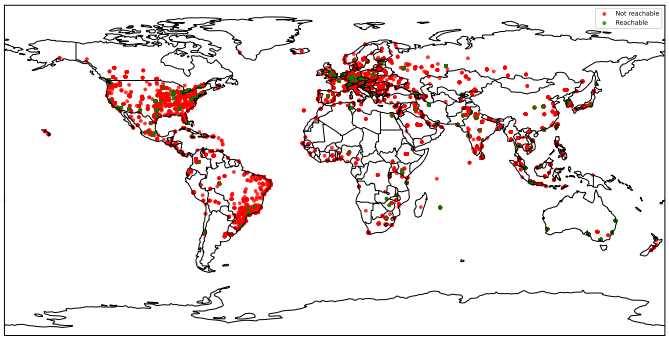

Fig. 5. Geographic distribution of DNS resolvers used by rDNSmt, colored by reachability. The figure shows that most DNS servers were not reachable.

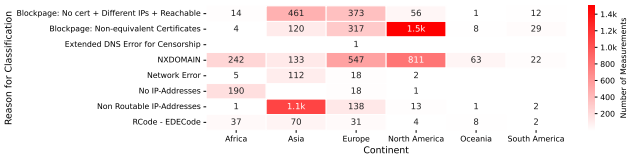

Fig. 6. Distribution of reasons for classification as "Censored" for reachable resolvers. EDE records are not used by DNS resolvers to indicate censorship. The most common behaviours for blocking domains are non-routable IP addresses, NXDOMAIN responses, or redirection to a block page.

as not censored due to matching IPs from test and control resolvers or resulted in equivalent TLS certificates.

The primary cause of censorship was inferred when the IPs resulted in different SSL certificates, suggesting the presence of a blockpage (see Fig. 6). Validating this prediction, using Censored Planet's signature-based patterns, showed that 1447 of the IPs led to confirmed blockpages and only 226 couldn't matched a pattern. This result highlights the reliability of this method for identifying blockpages when test and control resolvers return different IPs. Similarly important for censorship classification were cases in which the control resolver was able to correctly resolve the test domain, while the test resolver returned a valid but incorrect response *NXDOMAIN*, *No IP*, or a *Non-routable IP address*. These cases clearly indicate blocking by the resolver and do not require further analysis. The evaluation of blockpage prediction based on differing IPs from test and control resolvers, absence of TLS certificates, and presence of HTTP access shows limited reliability: 322 IPs matched known blockpage signatures, while no pattern was found for 1010 IPs. Only a single resolver returned an EDE code indicating blocking (code 15) for one domain. This demonstrates that EDE codes currently play only a minor role in identifying censorship through DNS resolvers.

Evaluation of rDNSmt Detected DNS Poisoning: If a resolver was marked as unreachable, a DNS request for the test domain was nevertheless sent to it to detect potential censorship devices. Such probing was only successful in China and Iran, with 1452 and 4337 blocked domains, respectively. In China, censorship devices responded with random, unreachable IP addresses, while in Iran, devices returned non-routable

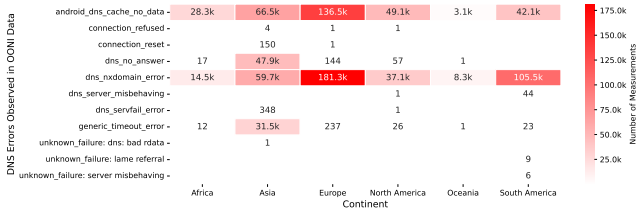

Fig. 7. Overview of OONI error responses leading to DNS anomalies. Some measurements resulted in "android_dns_cache_no_data", whose origin could not be determined. The most common error causing a DNS anomaly was NXDOMAIN.

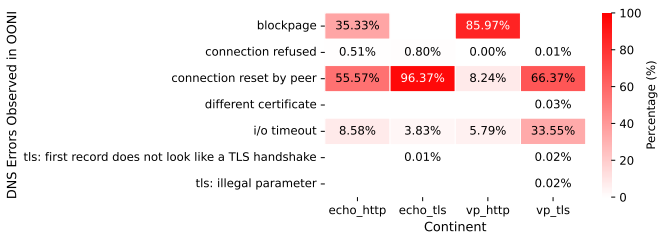

Fig. 8. Classification reasons for censored measurements using rVPmt. The most common technique for disrupting TLS connections was the injection of TCP RST packets, whereas HTTP connections were most frequently blocked through the injection of block pages.

IPs such as *10.10.34.34*, *10.10.34.35*, and *10.10.34.36*. Interestingly, for domains with *google* as the second-level domain, censorship devices in Iran responded with the correct IP addresses, possibly because these domains were previously blocked but are no longer subject to censorship.

Evaluation Including OONI's DNS Responses: When rDNSmt couldn't reach the test resolver, it used OONI's DNS responses. Measurements were classified as "not censored" primarily when the TLS certificates were equivalent or when the resolvers returned the same IP address. Many IPs returned by the test resolvers were unreachable, but this was less frequent with direct rDNSmt tests. The discrepancy likely arises from the time gap between OONI's measurements and rDNSmt's tests.

In the analysis of results classified as "censored", blockpage predictions, unroutable IP addresses, and error responses remained the primary indicators. The predictions showed slightly lower reliability compared to the direct rDNSmt tests: no blockpage pattern could be matched in $36\,\%$ of cases with mismatched TLS certificates and in $52\,\%$ of cases where the IP was only accessible via HTTP and differed from the control resolver's IPs. In cases where only the test resolver triggered an error response in the OONI measurements, nearly half were associated with the error type *android_dns_cache_no_data* (see Fig. 7). The exact origin of this error remains unclear. Therefore, these measurements were not considered as censored in future analyses. The remaining distribution of errors closely resembles the findings from the independent measurements. As in previous analyses, *NXDOMAIN*s are essential censorship patterns, alongside resolver timeouts and *dns_no_answer* (valid response was received with no IP in it). Other OONI error types occurred too infrequently to suggest a consistent censorship mechanism and were therefore treated as not censored in future classifications.

D. rVPmt

A total of 485815370 measurements were performed using rVPmt, with 2758320 classified as censored. For TLS-based connections, the predominant censorship mechanism is the injection of *connection reset* packets, as shown in Fig. 8. Additionally, the abrupt termination of connections, resulting in timeouts, also appears to be a common technique, although its overall share is lower. The remaining errors are interpreted as random occurrences and, due to their low frequency, will not be considered as evidence of censorship in the further analyses. In a total of 53 measurements, the TLS handshake to a VP resulted in a certificate matching the spoofed domain, although the VP responded with a different certificate for all other requests. For example, one of these VPs, located in Moscow, returned a different TLS certificate exclusively for the domain `cdn2.signal.org`, featuring a validity period of 10 years.

When analysing the censorship classifications made by rVPmt for HTTP requests to VPs (not those using the ECHO protocol), it is striking how often the result is a detected blockpage, over 700000 times. In comparison, TLS-based measurements resulted in fewer than 200000 censorship classifications, indicating a overvaluation. Validating the blockpage predictions using Censored Planet's patterns confirmed correct identification in $95\,\%$ of cases for Echo VPs, while accuracy dropped to $58\,\%$ for regular HTTP VPs. An analysis of these blockpages reveals two main challenges. First, some HTTP VPs respond too inconsistently to spoofed requests, making it difficult to establish a general approach for reliably determining whether a response originates from the VP itself. Second, certain corporate networks operating the VP deploy their own censorship devices, as evidenced by company-branded blockpages.

To extract blockpages originating from ISP-level or national censorship devices, the blockpage pattern sets provided by Censored Planet can once again be utilized. Their categorisation makes it possible to specifically isolate blockpages linked to national firewalls and ISP-level filtering. Across all protocols, it was observed that injecting TCP RST packets is a commonly used technique. In contrast, censorship mechanisms that silently terminate the connection, causing a timeout, were far less frequent. Blockpages continue to represent a significant share of censorship events in the case of HTTP traffic.

Analyzing the measurements that triggered the Russian national blockpage reveals that not all VPs involved are located in Russia. Some are based in Kyrgyzstan and, in six cases, Tajikistan (see Table 1). The most likely explanation is that Internet traffic from Kyrgyzstan and Tajikistan is partially routed through Russia, allowing the Russian censorship system to intercept and block these connections before they reach their

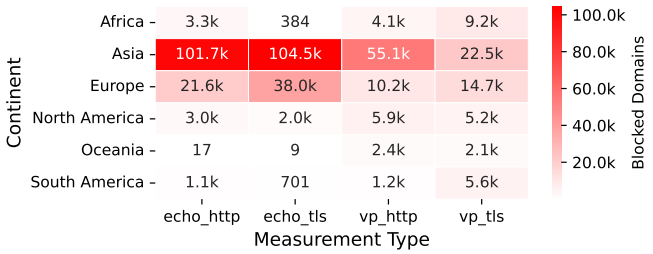

Fig. 9. Distinct blocked domains per continent based on VP remote measurements. The majority of blockages were observed in Asia, followed by Europe.

destination. This highlights a previously overlooked limitation of remote measurements, namely, that censorship may occur not only at the destination but also in transit countries along the route. At the same time, the data suggests that users in Kyrgyzstan and, to some extent, Tajikistan are subject to Russian censorship, as their traffic to the West is at least partially routed through Russian infrastructure. None of the other blockpages showed signs of incorrect allocation.

TABLE I
OVERVIEW OF MEASUREMENTS RESULTING IN THE RUSSIAN NATIONAL BLOCK PAGE, SORTED BY COUNTRY. IN ADDITION TO CONNECTIONS TO RUSSIA, CONNECTIONS TO KYRGYZSTAN AND TAJIKISTAN WERE ALSO DISRUPTED.

Country	Number of Distinct Domains
Russia	18277
Kyrgyzstan	3695
Tajikistan	6

The censored measurements, sorted by country and distinct domains, show that China, Iran, and Russia perform the highest number of blocks. However, the presence of the United States in the top 10 suggests a still notable false positive rate, mainly caused by *connection reset* packets. Since the majority of the top 35 countries are also classified as "Not Free" according to the Freedom on the Net Report [19], this indicates a generally functioning detection mechanism for censorship. The detected blocks using rVPmt are distributed across continents and VP types, as illustrated in Fig. 9.

E. The Final Blocklist

After combining the lists, as described in Section III, the aggregated results by country show the following pattern: Iran, Russia and China are at the top three countries, with the highest number of blocked domains. However, the continued presence of the United States among the top 10 suggests that some false positives remain. Overall, though, when comparing the results with the Freedom on the Net Report and manually checking selected entries, the approach gives a reasonable overall picture of global censorship. The applied methodology shows, in comparison to the initial situation of OONI anomalies, a clear improvement based on the expected censorship levels of each country if the freedom rating of the Freedom Report is used as a reference [19]. For example, instead of

Fig. 10. Percentage of RIPE Atlas measurements targeting domains found in the country's active blocklist (Top 35). Iran, Russia, and China show the highest potential risk, with more than 15% of the conducted measurements indicating possible blocking.

Brazil and France being among the top 3, Iran, China, and Russia now lead the ranking, with France dropping to 13th place. Except Oman, all countries that engage in domain blocking use DNS-based techniques. In addition to DNS-based blocking, several countries, including the United Arab Emirates, China, Egypt, India, Iran, Kyrgyzstan, Korea, Kazakhstan, Myanmar, Oman, Pakistan, the Russian Federation, Türkiye, the United States, and Vietnam, also show clear evidence of blocking based on the HTTP host header or the SNI. These methods require active interference by censorship devices, indicating a more advanced censorship infrastructure. Complete per-country statistics comparing advertised and measured blocklists are publicly available on WorldWideBlock.

If the blocklists were applied to the RIPE Atlas measurements to assess which measurements could potentially cause issues, either for the probe operators themselves or by introducing bias into the results, some measurements would appear problematic in certain countries (see Figure 10). For this analysis, the only base domains were compared to identify overlaps.

V. ETHICS

This research was conducted in accordance with established ethical guidelines for network measurement research. No third-party devices were used during the measurements. The VPs employed by rVPmt belong primarily to ISPs and government entities operating public-facing infrastructure, not to private individuals. Consequently, the risk of harm to VP hosts is minimal. The focus of this work is exclusively on identifying the presence of network blocks, not on circumventing them. Prohibited content was not retrieved at any time. To minimize risks to users in censored regions, identified blocked domains that are not already listed in published authoritative blocklists are not publicly disclosed. For RIPE Atlas, the analysis relied solely on publicly available measurement data; no additional probes were directed at sensitive domains. All measurements were conducted passively with respect to content access and do not engage in activities that would violate applicable laws or terms of service. This distinction ensures that the research stays within the bounds of ethical standards and respects legal frameworks

VI. CONCLUSION

The identified blocklists and their comparison with 2020 data clearly illustrate that states increasingly seek to assert digital sovereignty and regain control over online freedoms. Furthermore, it became evident that neither remote censorship measurements nor OONI data alone suffice to reliably attribute

censorship events to network operators. This is likely one reason why both Censored Planet and OONI do not (any more) label measurements as definitively censored (unless a blockpage is explicitly detected). However, by combining both approaches and focusing on overlapping signals, particularly in TLS- and HTTP-based blocking, censorship events can be more reliably interpreted. This allows an automated creation of blocklists with a significantly reduced risk of false positives. Nonetheless, because the methodology relies on OONI measurements and the availability of suitable VP within the same AS, completeness cannot be guaranteed and some false positives are likely still present. Especially since DNS servers were likely evaluated that implement other blocking policies, such as corporate filtering. To detect DNS blocking implemented by the resolver itself, remote measurements proved unsuitable.

Overall, domain blocking was observed in a total of 163 countries, with the majority of these measures being implemented without public transparency. With the exception of Lebanon, Indonesia, Kazakhstan, Kyrgyzstan, the United Kingdom and Switzerland, most discovered blocklists were associated with European countries Fig. 3 (RQ 1). With the exception of Oman, all countries that engage in domain blocking use DNS-based techniques. These typically include responding with NXDOMAIN, returning an unroutable IP address or redirecting to a block page (see Fig. 6 and Fig. 7). DNS poisoning could be actively confirmed only in China and Iran. Since most DNS resolver responses had to be taken from OONI measurements anyway, additional DNS poisoning results would already be included through that data. In addition to DNS-based blocking, there were also clear indications of blocking based on the HTTP host header or the SNI in numerous countries. These methods require active interference by censorship devices, indicating a more advanced censorship infrastructure. In most cases, a TCP RST packet was injected to terminate the connection. For blocks based on the HTTP host header, injecting a blockpage was also common (see Fig. 9) (RQ 3). The compiled blocklist identified and confirmed OONI anomalies by country and aims to reflect the current extent of blocking (RQ 2). The largest number of blocks was observed in China, Iran, and Russia. However, the relatively small values in some cases and the comparison with identified blocklists suggest that this does not represent the full extent of the actual blocklist in use and needs broader measurements by OONI around the world. In summary, the findings indicate that integrating multiple data sources enhances censorship detection accuracy. This approach underscores the need for collaboration, data sharing, and open-source tools to strengthen future research and global efforts against online censorship.

REFERENCES

[1] T. Berners-Lee, "Long live the web," *Scientific American*, vol. 303, no. 6, pp. 80–85, 2010.

[2] R. Sommese, A. Sperotto, A. Prado, J. {van der Ham}, and A. Affinito, "90th minute: A first look to collateral damages and efficacy of the italian piracy shield," in *2025 21th International Conference on Network and Service Management (CNSM)*. United States: IEEE, Aug. 2025, 21th International Conference on Network and Service Management, CNSM 2025, CNSM 2025 ; Conference date: 27-10-2025 Through 31-10-2025. [Online]. Available: https://www.cnsm-conf.org/2025/

[3] F. House, "Freedom on the net 2023," 2023, accessed: 2024-10-15. [Online]. Available: https://freedomhouse.org/sites/default/files/2023-10/Freedom-on-the-net-2023-DigitalBooklet.pdf

[4] V. Ververis, L. Lasota, T. Ermakova, and B. Fabian, "Website blocking in the european union: Network interference from the perspective of open internet," *Policy & Internet*, vol. 16, no. 1, pp. 121–148, 2024. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/poi3.367

[5] R. Barnes, A. Cooper, O. Kolkman, D. Thaler, and E. Nordmark, "Technical Considerations for Internet Service Blocking and Filtering," RFC 7754, Mar. 2016. [Online]. Available: https://www.rfc-editor.org/info/rfc7754

[6] E. Tsai, R. S. Raman, A. Prakash, and R. Ensafi, "Modeling and detecting internet censorship events," in *31st Annual Network and Distributed System Security Symposium, NDSS 2024, San Diego, California, USA, February 26 - March 1, 2024*. The Internet Society, 2024. [Online]. Available: https://www.ndss-symposium.org/ndss-paper/modeling-and-detecting-internet-censorship-events/

[7] A. Filastò and J. Appelbaum, "OONI: open observatory of network interference," in *2nd USENIX Workshop on Free and Open Communications on the Internet, FOCI '12, Bellevue, WA, USA, August 6, 2012*, R. Dingledine and J. Wright, Eds. USENIX Association, 2012. [Online]. Available: https://www.usenix.org/conference/foci12/workshop-program/presentation/filast\%C3\%B2

[8] Open Observatory of Network Interference (OONI), "Web connectivity measurement: 4667fb6a2ac033ff," 2025, accessed on 2025-05-25. [Online]. Available: https://explorer.ooni.org/m/20250506065010.654013_DE_webconnectivity_4667fb6a2ac033ff

[9] R. S. Raman, P. Shenoy, K. Kohls, and R. Ensafi, "Censored planet: An internet-wide, longitudinal censorship observatory," in *CCS '20: 2020 ACM SIGSAC Conference on Computer and Communications Security, Virtual Event, USA, November 9-13, 2020*, J. Ligatti, X. Ou, J. Katz, and G. Vigna, Eds. ACM, 2020, pp. 49–66. [Online]. Available: https://doi.org/10.1145/3372297.3417883

[10] OONI, "New ooni explorer features!" 2023, accessed: 2024-11-15. [Online]. Available: https://ooni.org/post/2023-new-explorer-features/

[11] R. S. Raman, P. Shenoy, K. Kohls, and R. Ensafi, "Censored planet: An internet-wide, longitudinal censorship observatory," in *CCS '20: 2020 ACM SIGSAC Conference on Computer and Communications Security, Virtual Event, USA, November 9-13, 2020*, J. Ligatti, X. Ou, J. Katz, and G. Vigna, Eds. ACM, 2020, pp. 49–66. [Online]. Available: https://doi.org/10.1145/3372297.3417883

[12] A. Vyas, R. S. Raman, N. Ceccio, P. M. Lutscher, and R. Ensafi, "Lost in transmission: Investigating filtering of COVID-19 websites," in *Financial Cryptography and Data Security - 25th International Conference, FC 2021*, N. Borisov and C. Díaz, Eds., vol. 12675. Springer, 2021, pp. 417–436. [Online]. Available: https://doi.org/10.1007/978-3-662-64331-0_22

[13] S. Fujii and T. Sato, "10 pdnss 10 colors: A measurement study of protective dns," *Journal of Information Processing*, vol. 33, pp. 608–618, 2025.

[14] R. S. Raman, M. Wang, J. Dalek, J. R. Mayer, and R. Ensafi, "Network measurement methods for locating and examining censorship devices," in *CoNEXT 2022*. ACM, 2022, pp. 18–34. [Online]. Available: https://doi.org/10.1145/3555050.3569133

[15] S. Nourin, V. H. Tran, X. Jiang, K. Bock, N. Feamster, N. P. Hoang, and D. Levin, "Measuring and evading turkmenistan's internet censorship: A case study in large-scale measurements of a low-penetration country," in *Proceedings of the ACM Web Conference 2023, WWW 2023*. ACM, 2023, pp. 1969–1979. [Online]. Available: https://doi.org/10.1145/3543507.3583189

[16] DeepL SE, "Deepl translator," https://www.deepl.com/translator, 2025, accessed: 2025-09-18.

[17] Censored Planet, "assets-censoredplanet," https://assets.censoredplanet.org/, accessed: 2026-02-28.

[18] R. Kisteleki. (2016) Ethics of ripe atlas measurements. Accessed on 2025-05-14. [Online]. Available: https://labs.ripe.net/author/kistel/ethics-of-ripe-atlas-measurements/

[19] Freedom House, "Freedom in the World 2025: Country Scores," 2025, accessed on 2025-05-14. [Online]. Available: https://freedomhouse.org/countries/freedom-world/scores