

Towards a Comprehensive Understanding of Russian Transit Censorship

Aaron Ortwein
University of Maryland

Kevin Bock
University of Maryland

Dave Levin
University of Maryland

Abstract

Within the past decade, Russia has significantly increased its censorship efforts. This raises the question as to whether Russian ISPs are applying their censorship policies not only to traffic terminating in Russia, but also to traffic that simply transits Russia. Evidence of this “collateral damage” in central Asia due to filtering by upstream Russian network providers has been noticed in previous research, but the full extent of it has yet to be studied.

In this work, we present first steps toward a comprehensive study of the collateral damage of Russian censorship. We scan the IP address spaces of 18 countries surrounding Russia while attempting to elicit responses from Russian censors. We identify Russian collateral damage affects at least some of the traffic for 8 of these 18 countries, and that at least 6 ASes are responsible for censorship of transit traffic. Our results highlight the need for further study of collateral damage globally.

1 Introduction

Censoring nation-states analyze and tamper with traffic that does not adhere to their national policies. While typically thought of as a way to control information into or out of their own country, these policies can sometimes be applied—intentionally or not—to traffic that merely transits *through* their country. We refer to this as *transit censorship*; it is often considered a form of collateral damage, in that it results in ostensibly unintentional over-blocking.

Previous research has observed Russia blocking transit traffic destined for Kazakhstan, Kyrgyzstan, and Uzbekistan [16, 20]. Additionally, users in countries such as Armenia, Azerbaijan, Belarus, Moldova, and Ukraine have regularly reported receiving Russian blockpages [23, 24]. Russian transit censorship has occasionally reached even further than Russia’s neighbors, such as during a 2016 incident in which DDoS mitigations deployed by rutracker.org temporarily caused some international traffic to be routed through AS20485 (Joint Stock Company Transtelecom) [16].

However, these prior observations have largely been anecdotal; the full extent of Russian transit censorship has yet to be studied. There are two reasons to believe that Russian transit censorship is more common than previously reported. *First*, growth of Russia’s information controls over the past decade has been rapid and even haphazard at times. In 2018, Russia blocked Telegram by blocking millions of IP addresses belonging to Amazon and Google’s cloud hosting platforms, additionally blocking many unrelated websites [26]. In 2021, Russia’s throttling devices initially incorrectly implemented regular expression matching for the Twitter link shortener domain t.co, throttling connections to any domain containing the string “t.co” [14]. Domains ending in “twitter.com” were also throttled briefly [29].

Second, the architecture of the Internet and Russia’s censorship infrastructure provide the opportunity for inflicting collateral damage upon other countries. Acharya et al. found that a small set of ASes, some of which are located in censoring regimes such as Russia, appear in 90% of routing paths to popular websites, and that approximately 11% of routing paths to popular websites transit Russia [1]. Some ISPs in Russia rely on upstream providers to perform censorship [28]. If the Russian ASes responsible for routing multinational Internet traffic coincide with those responsible for censorship—and if censorship devices operate over all traffic regardless of its source address—then that could lead to a significant amount of transit censorship.

In this work, we take a first step towards measuring the extent of Russian transit censorship. We find that the resulting collateral damage is more widespread than observed in previous research, affecting at least 8 countries surrounding Russia and attributable to at least 6 ASes. Due to multiple limitations of our measurements, it is important to note that these results are preliminary. We hope to further investigate the reach of Russia’s and other censorship infrastructures in future work.

2 Background

Censorship Models Some countries such as China and Iran have centralized censorship models in which all Internet traffic passes through and is subject to censorship at one of very few state-operated points of control [4, 15, 27]. However, governments that wish to enforce information controls and lack unilateral control over networks often opt for a decentralized censorship model in which legal institutions dictate which resources should be blocked and when, but the responsibility of implementing the technical mechanisms for blocking falls upon individual ISPs. As a result, the blocking methods, targets, and efficacy can vary between ASes and even their constituent networks.

Since the 2012 inception of a national blocklist [22] maintained by Russia’s information controls authority, Roskomnadzor, Russia has built a primarily decentralized censorship system. Data center networks tend to favor blocking traffic destined for blocklisted IP addresses, while residential ISPs often return blockpages via DNS manipulation or injection into connections containing forbidden keywords found by Deep Packet Inspection (DPI) of application-layer headers and payloads [21]. However, with the signing of the 2019 Sovereign RuNet law, Russian ISPs were required to install TSPU devices, DPI systems centrally controlled by Roskomnadzor, into their networks [9]. The TSPU devices have recently been used to throttle Twitter and block access to resources related to the Russian war effort in Ukraine uniformly across ISPs [28, 29]. The recent use of TSPU devices marks a shift towards a centralized censorship model built over a decentralized infrastructure, although the individual ISP censorship systems continue to operate [20, 28].

Collateral Damage Collateral damage refers to the (typically inadvertent) over-blocking of Internet resources. Traditionally, collateral damage occurs when a censoring regime implements blocking such that its own citizens are unable to access significantly more webpages or websites than the blocking target. The prevalence of content delivery networks (CDNs), which host many websites behind a few IP addresses, has diminished the popularity of IP-based blocking due to the risk of additionally denying access to many other websites hosted at the same IP address as a blocking target [25]. Even with the finer-grained blocking policies enabled by Deep Packet Inspection of application-layer headers and payloads, relaxed regular expression-based blocking rules may match many domains that are unrelated to the blocking target [29].

However, collateral damage can also affect Internet users from other countries when traffic from one country to another is subject to filtering by a third, censoring country while in transit. In this case, censorship mechanisms block more users than perhaps intended, or otherwise block users who may not expect to be blocked. Our work focuses on identifying this form of collateral damage by Russia’s censorship apparatus.

3 Related Work

Previous research has investigated the blocking of foreign Internet traffic that traverses a censored network. Acharya et al. mapped Internet routes to popular websites and estimated the potential collateral damage from routes through censoring countries such as China, India, and Russia, but they assume that all transit links in these countries are censored and did not measure the actual extent of collateral damage [1]. Anonymous work has quantified the extent of collateral damage by China’s DNS injectors on DNS requests destined for root server IP addresses in China [3]. Similarly, the Citizen Lab reported on a number of URLs visited from Oman subject to upstream filtering in India [8]. Cho et al. localized censors at a global scale and identified 18 ASes that censor transit traffic; they explicitly reported only the most commonly observed transit censors, which were located in Hong Kong, Sweden, and Japan [7].

Censorship of transit traffic has also been observed in Russia. A 2016 study by Ukrainian ISP NetAssist LLC found evidence that select networks in Kazakhstan, Kyrgyzstan, and Uzbekistan were impacted by censorship of transit traffic through AS3216 (PJSC Vimpelcom) [16]. NetAssist speculated that Armenia, Azerbaijan, and Georgia might be affected by Russian censorship, but that no European countries were. Additionally, recent work by Raman et al. discovered that over one-third of their remote measurements to Kazakhstan timed out in one of two ASes in Russia: AS31133 (PJSC MegaFon) and AS43727 (JSC Kvant-Telekom) [20]. However, the true extent of the collateral damage is still unknown.

4 Methodology

4.1 Gathering Blockpages

We limit the scope of this work to transit censorship in the form of injected blockpage responses—rather than blocking via packet drops or TCP RSTs—because blockpages are immediately attributable to interference by Russian ISPs. Indeed, many Russian blockpages tend to cite federal law and link to the national blocklist, distinguishing them from other blockpages and server error pages [19, 21]. Detecting transit censorship in the form of dropped or reset connections requires distinguishing these behaviors from transient network errors and determining where on the routing path they occur.

The variability in blocking methods, targets, and efficacy arising from Russia’s decentralized censorship infrastructure presents a challenge for selecting a domain that will elicit censorship responses from all ISP censorship systems. We use ZGrab [10] to complete the TCP handshake with and send an HTTP GET request containing a forbidden Host header to each host in the Russian IP address space. We identify Russian blockpages by matching HTTP responses against regular expressions developed by OONI [18] and Censored

Planet [6, 19]. We then use the pyasn [12] Python module to map each destination IP address for which we received a blockpage to its ASN. We choose facebook.com, finding that forbidden requests destined to over 900 ASes elicited a blockpage response.

4.2 Scanning for Transit Censorship

We use a version of ZMap [11] extended by Bock et al. [5] to scan the entire IP address spaces of 18 countries surrounding Russia, primarily in eastern Europe and central Asia. IP address spaces are obtained from NRO Extended Allocation and Assignment Reports [17], which publish IP range and ASN allocations on a daily basis. We note that the countries associated with allocations correspond to the location of the organization, though the physical networks may be located elsewhere.

We send a SYN packet with sequence number s followed by a PSH+ACK packet with sequence number $s + 1$ and a HTTP GET request whose Host header contains facebook.com. Bock et al. found that this packet sequence is effective in eliciting responses from middleboxes [5]. Censors generally do not expect to see all packets in a connection due to routing path variance and route asymmetry, where packets follow a different path to the destination than when returning from the destination. A SYN followed by a PSH+ACK looks like a standard TCP connection without the server’s SYN+ACK and the client’s ACK, so many censors operate under the assumption that they missed part of the connection. Because we do not complete the TCP handshake—our PSH+ACK packet does not increment the acknowledgement number returned by a server’s SYN+ACK—we expect that responses with a payload are sent by middleboxes, though not necessarily Russian censors. Since our measurement strategy does not require a server response, we can measure a much wider range of destination IP addresses victim to transit censorship by a censor. However, a core limitation is that we do not observe transit censorship by censors that must observe a complete TCP handshake before blocking. To account for differences in censorship due to routing path variance, we perform scanning from three vantage points: a research machine run by a public US university, as well as two AWS instances located in Sydney and Tokyo.

In order to save space, the ZMap scan module developed by Bock et al. does not record the actual response packets, instead representing them as a tuple including the source IP address, packet size, payload length, and TCP flags. For each distinct (packet size, payload length, TCP flags) triplet with a nonzero payload length, we selected a random IP address from amongst those that sent such a response and issued a HTTP GET request with facebook.com in the Host header. If the HTTP response matched one of the OONI or Censored Planet Russian blockpage regular expressions, we treated all instances of the (packet size, payload length, TCP flags) tu-




Figure 1: Countries affected by Russian transit censorship (yellow), and our measurement vantage points (blue).

ple as a blockpage response. We then counted the number of unique source IP addresses which returned a blockpage, yielding the number of IP addresses affected by transit censorship from the perspective of the vantage point.

4.3 Localizing Censorship Devices

We determine which ASes are responsible for transit censorship by sending TTL-limited forbidden GET requests, each preceded by an initial SYN packet, to a randomly-selected IP address amongst those in the same aggregated network prefix and that share the same blockpage response fingerprint. We compare the minimum TTL value t that consistently elicits a Russian blockpage response to traceroutes to the destination IP. If traceroute can identify the IP address at hop t , we find its AS via a WHOIS lookup.

However, this strategy is complicated by multiple factors. First, hop t of the traceroute may differ from hop t of the forbidden request due to routing path variance. Additionally, for multiple ASes, we observe a phenomenon described in recent work [13, 20]: censorship devices copy the TTL value of a forbidden packet into their response, so the TTL of the request must be at least twice the number of hops to the censor in order for the client to receive the blockpage. Route asymmetry may cause the outbound and return path lengths to differ, so hop $\frac{t}{2}$ of the traceroute may not belong to the censoring AS. Finally, some censors hide from traceroute by refusing to either respond or decrement the TTL of forwarded packets. We therefore also perform lookups on nearby hops and manually check whether the ISP is identified by the blockpage URL or content. Otherwise, we may misidentify the AS responsible for transit censorship.

5 Preliminary Results

Where do we observe transit censorship? From our ZMap scans, we observe Russian censorship responses to traffic en route to at least one network in each of Afghanistan, Azerbaijan, Kyrgyzstan, Kazakhstan, Lithuania, South Korea, Tajik-

AS	Organization
AS3216	PJSC Vimpelcom
AS25227	JSC Avantel
AS35816	Lancom Ltd
AS47203	JSC CrimeaTelecom
AS60299	Mezhdugorodnyaya Mezhdunarodnaya Telefonnaya Stanciya Ltd
AS201776	Miranda-Media Ltd

Table 1: Russian ISPs responsible for transit censorship.

istan, and Ukraine. Among these 8, only 2 had been reported (Kyrgyzstan and Kazakhstan) in previous studies. Unlike previous studies, we do not observe transit censorship affecting Uzbekistan.

The countries for which we observe transit censorship and the blockpages encountered differ between our vantage points due to differences in routing paths to the destination IP addresses. Our US-based vantage point experienced the most diverse range of censorship in terms of the number of impacted countries: ZMap scans detected blockpage fingerprints in all 8 countries. ZMap scans from both our Sydney and Tokyo vantage points only observed transit censorship affecting Kazakhstan and Ukraine.

Which ASes are responsible for transit censorship? We find that filtering of transit traffic occurs in at least 6 ASes, which are listed in Table 1. Of these, 4 return blockpages identifying the censoring ISP. AS60299 (Mezhdugorodnyaya Mezhdunarodnaya Telefonnaya Stanciya Ltd) and AS201776 (Miranda-Media Ltd) deploy commercial DPI technology manufactured by Russian company VAS Experts.

AS3216 (PJSC Vimpelcom) has the furthest reach in terms of number of countries affected, delivering our US-based client blockpages for traffic destined to certain IP addresses in Afghanistan, Azerbaijan, Kyrgyzstan, Kazakhstan, Lithuania, South Korea, Tajikistan, and Ukraine. However, the number of unique destination IP addresses for which we receive a blockpage is relatively low, regardless of our vantage point. From all three of our vantage points, no country experiences transit censorship by AS3216 (PJSC Vimpelcom) for more than 1,000 IP addresses. Moreover, our Sydney vantage point only observes transit censorship by AS3216 (PJSC Vimpelcom) in Afghanistan, Kyrgyzstan, and Kazakhstan, and our Tokyo vantage point does not observe any transit censorship by AS3216 (PJSC Vimpelcom).

In terms of the number of destination IP addresses for which we experience transit censorship, AS201776 (Miranda-Media Ltd) is responsible for most blocking by far, but its impact is limited to a single country. Both our US and Sydney vantage points observed transit censorship by AS201776 (Miranda-Media Ltd) for approximately 16,000 IP addresses in Ukraine. Our Tokyo vantage point observes significantly

fewer IP addresses impacted by these ASes. We hypothesize that packet loss arising from the scanning rate and differences in vantage point resources is a contributing factor to this result.

Overall, Ukraine is subject to transit censorship by the most Russian ASes, likely due to recent re-routing of Ukrainian Internet traffic through the Russian telecommunications infrastructure [2]. In addition to that by AS3216 (PJSC Vimpelcom) and AS201776 (Miranda-Media Ltd), we observe transit censorship by AS25227 (JSC Avantel), AS35816 (Lancom Ltd), and AS47203 (JSC CrimeaTelecom). From our US and Sydney vantage points, AS25227 (JSC Avantel) impacts routes to over 1,500 IP addresses. The transit censorship by the latter two ASes is relatively small: AS47203 (JSC CrimeaTelecom) impacts just under 300 IP addresses from our US and Sydney vantage points, and AS35816 (Lancom Ltd) impacts just a single IP address. We also observe one blockpage from only our Tokyo vantage point that affects nearly 300 IP addresses but we are currently unable to attribute to an AS via traceroutes alone; the blockpage is hosted in AS6789 (CRELCOM LLC), but traffic does not appear to pass through this network, and the minimum TTL value that consistently triggers censorship is twice the hop distance of the server itself.

6 Conclusion

We demonstrate that the collateral damage of Russia’s censorship infrastructure is more extensive than previously known. We emphasize, however, that these are preliminary results. In particular, our initial study has two key limitations: we constrained ourselves to only look at blockpages, and our few vantage points have low coverage of all possible routing paths through Russian networks. As a result of these limitations, we anticipate that our findings are a lower bound of the true impact of Russian censorship. In our future work, we plan to conduct a more comprehensive study with many globally-distributed vantage points, and to detect other forms of censorship beyond blockpages.

Transit censorship appears to be a relatively understudied aspect of nation-state censorship. We hope that our findings inspire the broader community to more fully include transit censorship in their measurements. To this end, in the future, we hope to expand our efforts to perform a global measurement of the collateral damage caused by transit censorship.

Acknowledgments

We thank the anonymous reviewers and our shepherd, Arturo Filastò, for their helpful feedback. This work was supported in part by NSF award CNS-1943240.

References

- [1] H. B. Acharya, Sambuddho Chakravarty, and Devashish Gosain. Few Throats to Choke: On the Current Structure of the Internet. In *IEEE Conference on Local Computer Networks (LCN)*, 2017.
- [2] Adam Satariano. How Russia Took Over Ukraine’s Internet in Occupied Territories. <https://www.nytimes.com/interactive/2022/08/09/technology/ukraine-internet-russia-censorship.html>, August 2022.
- [3] Anonymous. The collateral damage of internet censorship by dns injection. *SIGCOMM Comput. Commun. Rev.*, 2012.
- [4] Simurgh Aryan, Homa Aryan, and J. Alex Halderman. Internet censorship in Iran: A first look. In *USENIX Workshop on Free and Open Communications on the Internet (FOCI)*, 2013.
- [5] Kevin Bock, Abdulrahman Alaraj, Yair Fax, Kyle Hurley, Eric Wustrow, and Dave Levin. Weaponizing Middleboxes for TCP Reflected Amplification. In *USENIX Security Symposium*, 2021.
- [6] CensoredPlanet. Censored Planet assets. <https://assets.censoredplanet.org/>.
- [7] Shinyoung Cho, Rishab Nithyanand, Abbas Razaghpanah, and Phillipa Gill. A churn for the better: Localizing censorship using network-level path churn and network tomography. In *ACM Conference on emerging Networking EXperiments and Technologies (CoNEXT)*, 2017.
- [8] Jakub Dalek, Ron Deibert, Masashi Crete-Nishihata, Adam Senft, Helmi Noman, and Greg Wiseman. Routing Gone Wild: Documenting upstream filtering in Oman via India. <https://citizenlab.ca/2012/07/routing-gone-wild/>, July 2013.
- [9] Digital Russia. President Signs Sustainable Runet Act. <https://d-russia.ru/prezident-podpisa-zakon-ob-ustojchivom-runete.html>, May 2019.
- [10] Zakir Durumeric, David Adrian, Ariana Mirian, Michael Bailey, and J. Alex Halderman. A search engine backed by internet-wide scanning. In *ACM Conference on Computer and Communications Security (CCS)*, 2015.
- [11] Zakir Durumeric, Eric Wustrow, and J. Alex Halderman. ZMap: Fast Internet-wide Scanning and its Security Applications. In *USENIX Security Symposium*, 2013.
- [12] Hadi Asghari and Arman Noroozian. pyasn. <https://pypi.org/project/pyasn/>, 2020.
- [13] Lin Jin, Shuai Hao, Haining Wang, and Chase Cotton. Understanding the practices of global censorship through accurate, end-to-end measurements. *Proc. ACM Meas. Anal. Comput. Syst.*, 2021.
- [14] libneko. Twitter slowdown in Russia. <https://ntc.party/t/twitter/907>.
- [15] Bill Marczak, Nicholas Weaver, Jakub Dalek, Roya Ensafi, David Fifield, Sarah McKune, Arn Rey, John Scott-Railton, Ron Deibert, and Vern Paxson. An Analysis of China’s “Great Cannon”. In *USENIX Workshop on Free and Open Communications on the Internet (FOCI)*, 2015.
- [16] NetAssist LLC. Research of nationwide blacklist censorship effect on customers Internet access in nearby countries. https://ripe72.ripe.net/presentations/76-russian_censorship2.pdf, 2016.
- [17] NRO. Extended allocation and assignment reports. <https://www.nro.net/about/rirs/statistics/>.
- [18] OONI. Blocking Fingerprints. <https://github.com/ooni/blocking-fingerprints>.
- [19] Ram Sundara Raman, Adrian Stoll, Jakub Dalek, Armin Sarabi, Reethika Ramesh, Will Scott, and Roya Ensafi. Measuring the Deployment of Network Censorship Filters at Global Scale. In *Network and Distributed System Security Symposium (NDSS)*, 2020.
- [20] Ram Sundara Raman, Mona Wang, Jakub Dalek, Jonathan Mayer, and Roya Ensafi. Network measurement methods for locating and examining censorship devices. In *ACM Conference on emerging Networking EXperiments and Technologies (CoNEXT)*, 2022.
- [21] Reethika Ramesh, Ram Sundara Raman, Matthew Bernhard, Victor Ongkowitz, Leonid Evdokimov, Anne Edmundson, Steven Sprecher, Muhammad Ikram, and Roya Ensafi. Decentralized control: A case study of russia. In *Network and Distributed System Security Symposium (NDSS)*, 2020.
- [22] Registry of Banned Sites. <https://blocklist.rkn.gov.ru/>.
- [23] Roskomsvoboda. Rostelecom removes two-year blockage of full-fledged Internet access for subscribers from other countries. <https://roskomsvoboda.org/8690/>, September 2014.
- [24] Roskomsvoboda. Beeline blocks sites banned in Russia and outside the country. <https://roskomsvoboda.org/post/bilajn-blokiruet-zapreschyonnyie-v-rossi/>, September 2016.

- [25] Benjamin VanderSloot, Allison McDonald, Will Scott, J. Alex Halderman, and Roya Ensafi. Quack: Scalable Remote Measurement of Application-Layer Censorship. In *USENIX Security Symposium*, 2018.
- [26] Vlad Savov. Russia’s Telegram ban is a big, convoluted mess. <https://www.theverge.com/2018/4/17/17246150/telegram-russia-ban>, April 2018.
- [27] Xueyang Xu, Zhuoqing Mao, and J. Halderman. Internet censorship in china: Where does the filtering occur? In *Passive and Active Network Measurement Conference (PAM)*, 2011.
- [28] Diwen Xue, Benjamin Mixon-Baca, ValdikSS, Anna Ablove, Beau Kujath, Jedidiah R. Crandall, and Roya Ensafi. TSPU: Russia’s Decentralized Censorship System. In *ACM Internet Measurement Conference (IMC)*, 2022.
- [29] Diwen Xue, Reethika Ramesh, ValdikSS, Leonid Evdokimov, Andrey Viktorov, Arham Jain, Eric Wustrow, Simone Basso, and Roya Ensafi. Throttling Twitter: An Emerging Censorship Technique in Russia. In *ACM Internet Measurement Conference (IMC)*, 2021.