11-28-2017

# A Churn for the Better: Localizing Censorship Using Network-Level Path Churn and Network Tomography

Shinyoung Cho
*Stony Brook University*, scho@smith.edu

Abbas Razaghpanah
*Stony Brook University*

Rishab Nithyanand
*Data and Society Research Institute*

Phillipa Gill
*University of Massachusetts Amherst*

# A Churn for the Better

## Localizing Censorship using Network-level Path Churn and Network Tomography

Shinyoung Cho [1], Rishab Nithyanand [1], Abbas Razaghpanah [1], Phillipa Gill [2]

[1] Stony Brook University, [2] University of Massachusetts, Amherst



## ABSTRACT

Recent years have seen the Internet become a key vehicle for citizens around the globe to express political opinions and organize protests. This fact has not gone unnoticed, with countries around the world repurposing network management tools (*e.g.,* URL filtering products) and protocols (*e.g.,* BGP, DNS) for censorship. However, repurposing these products can have unintended *international* impact, which we refer to as "censorship leakage". While there have been anecdotal reports of censorship leakage, there has yet to be a systematic study of censorship leakage at a global scale.

In this paper, we combine a global censorship measurement platform (ICLab) with a general-purpose technique – *boolean network tomography* – to identify which AS on a network path is performing censorship. At a high-level, our approach exploits BGP churn to narrow down the set of potential censoring ASes by over 95%. We exactly identify 65 censoring ASes and find that the anomalies introduced by 24 of the 65 censoring ASes have an impact on users located in regions outside the jurisdiction of the censoring AS, resulting in the *leaking* of regional censorship policies.

## 1 INTRODUCTION

The Internet is now regarded as part of the critical infrastructure, with citizens relying on it for dissemination of information and organizing political action. Consequently, governments and network-level entities – *e.g.,* Autonomous Systems (ASes) – are implementing forms of censorship to restrict access to specific content, in many cases, repurposing existing network management tools [10] and protocols (*e.g.,* DNS [1, 22], BGP [4]) to filter Internet content. The past decade has seen numerous instances, where this repurposing of protocols has had unintended international impact, which we refer to as "censorship leakage." These have included the 2008 hijack of YouTube traffic by Pakistan Telecom [4], and cases of DNS root-servers located in China impacting international users [3]. While specific instances of censorship leakage have been a boon for researchers (in particular those studying China) [1, 9, 22, 34, 36], we lack a broader understanding of the prevalence of this phenomenon. Part of the challenge of understanding censorship leakage on a global scale is developing a technique that is able to identify censorship leakage in general vs. looking for specific instances (*e.g.,* injected RST or DNS packets from China [1, 34]). This challenge is compounded by a general lack of vantage points that are available for measuring censorship on an ongoing basis. We address these challenges by combining the ICLab measurement platform with the idea of *boolean network tomography* [33]. ICLab is a platform of ~1,000 globally distributed vantage point that has been performing measurements of censorship on an ongoing basis since November 2015 (more details in §2.1).

Our intuition is that we can observe multiple tests from a given vantage point to a given destination and that, if there is sufficient path churn between the vantage point and destination, we can create a set of boolean constraints where the constraint is true if censorship is observed, and false otherwise. In the case where censorship is observed, it must be the case that at least one autonomous system (AS) on the path is performing censorship. We can then input these constraints into an off-the-shelf SAT solver to identify the AS performing censorship. In this study, we demonstrate the applicability of this intuition by answering the following key questions: (1) Is there enough path churn observed in our measurements to create a solvable set of constraints? We need to validate that we have enough variability in paths, especially in the cases where we observe censorship, to create a solvable set of constraints to narrow down the set of potential censors. (2) Will our constraints generate a small set of potential censoring ASes? We want to make sure that the set of potential censoring ASes is not intractably large, making it impossible to exactly identify ASes responsible for implementing censorship. While answering these questions, we make the following contributions:

**Problem reformulation.** We demonstrate how measurements gathered by the ICLab platform can be used to formulate a boolean network tomography problem solvable by off-the-shelf SAT solvers. Our approach carries over to other measurement databases such as those generated by the OONI [16] and the M-Lab [23] platforms.

**Measuring and exploiting network-level churn.** We show that the instability of network-level paths can act as a substitute for strategically placed internal monitors. Specifically, we show that 25%, 30%, 28%, and 67% of paths between ICLab vantage points and web servers are observed to change over periods of one day, week, month, and year.

Email: {shicho, rnithyanand, arazaghpanah}@cs.stonybrook.edu, phillipa@cs.umass.edu.

These changes are found to significantly improve the solvability of our constructed SAT problems.

**Identifying censors and censorship leakage.** We empirically demonstrate that our approach allows us to reduce the size of the set of potential censoring ASes by over 95%, on average. Further, we exactly identify 65 censoring ASes located in 30 different countries. Our study also identifies *leakage* of censorship policies – *i.e.,* cases where censoring ASes blocked access to content even for users outside their country of operation. Specifically, we find that 32 and 24 of the censoring ASes leak censorship to other ASes and countries, respectively.

## 2 BACKGROUND & RELATED WORK

**Censorship measurement.** Much of previous work has focused on understanding how censorship is performed by network-level entities. Studies have shown that censors may restrict access to content by injecting incorrect DNS replies [1, 26], sending TCP reset packets spuriously [2, 21], using off-the-shelf filtering and blocking tools [11], or throttling connections to censored content [2]. Our study builds off of related work on large-scale longitudinal censorship measurement systems. Specifically, we use data collected by the ICLab platform [29] to detect censorship and generate constraints. Conceptually, our techniques could be applied to other platforms such as OONI [16] as well.

**Fault localization.** Other studies have approached the problem of network failure localization with different perspectives. Lifeguard [20] relies on historical control-plane measurements and active probing to automatically identify and route around network failures via crafted BGP messages. Feamster *et al.*[13] measure the effectiveness of reactive routing around node failures. Their approach localizes failures in real-time by analyzing results of active probes, including pings and traceroutes, between vantage points. While other work in the area has focused on identifying the root cause of path changes on the Internet [14, 18, 28, 31, 35], we focus on localizing faults and errors that do trigger path changes.

**Boolean network tomography.** Network tomography [33] typically involves using end-to-end measurements and a set of monitors within the network to uncover hidden node values (which in the case of boolean network tomography, may only take the values True or False). Monitors are used to ensure that appropriate end-to-end measurements may be performed to unveil specific node characteristics. We use the data gathered by the ICLab platform as end-to-end measurements from which we identify nodes (ASes) implementing specific types of censorship. Unlike typical boolean network tomography problems, our study is limited by the absence of strategically located monitors from which end-to-end measurements can be gathered. However, we show that due to the churn of network-level paths, we are still able to use boolean network tomography to identify censoring ASes.

Several studies have focused on the problem of error localization through boolean network tomography. Ma *et al.*[24] focus on identifying the conditions, monitor locations, and probing mechanisms that are necessary for fault localization through boolean network tomography. Dhamdhere *et al.*[12] use a boolean network tomography approach in conjunction with "troubleshooting sensors" (monitors) located within the network to identify misconfigured routers responsible for network failures. Other work has applied boolean network tomography to identify areas of heavy congestion [32] and packet loss [7, 8]. In this paper, we use boolean network tomography combined with the network-level path churn in routing protocols (as a substitute for specific monitors) to identify ASes that introduce censorship related anomalies.

### 2.1 The ICLab Dataset

We rely on data gathered by the ICLab censorship measurement platform [29] as a source for end-to-end measurements. The ICLab platform repetitively performs a variety of measurements between a set of over 1K globally distributed vantage points and web-servers hosting regionally sensitive content. The platform aims to (1) identify content being censored, (2) understand how censorship is implemented, and (3) record changes in censorship policies over time. Specifically, ICLab identifies the following anomalies as indicative of potential censorship:

**DNS anomalies.** DNS anomalies occur when a censor injects DNS responses to queries issued by a client. The idea being that if the censor is closer to the client, their injected packet will arrive at the client before the response from the DNS resolver. DNS injection can be detected by observing two response packets for the same DNS request. The ICLab platform identifies DNS injection by making DNS queries for domain names using both, Google's DNS resolver (`8.8.8.8`) and the default resolver of the vantage point, and then observing the number of response packets received, in each case. If a second DNS response packet is received within two seconds of the first, an anomaly is reported.

**SEQNO and TTL anomalies.** A packet injector will often have attributes that differ from the legitimate server for a connection. ICLab platform issues HTTP GET requests and records all responses (while following redirects). The platform then analyzes raw packet captures to identify anomalies. Specifically, we compare the IP TTL header on the SYNACK packet of the connection with subsequent packets. This relies on the assumption that a censor will not be fast enough to act prior to the SYNACK being sent by the server. Further, injected packets will often not be able to perfectly mimic the TCP state of the server [34]. We look for cases where there are overlapping sequence numbers between packets or gaps in sequence numbers. These sequence number anomalies, especially when combined with packets having the RST flag set (to close the connection) are likely indicators of censorship.

| Period | 2016-05 ~ 2017-05 |
|---|---|
| Unique URLs | 774 |
| AS Vantage Points | 539 |
| Destination ASes | 620 |
| Countries | 219 |
| Measurements | 4.9M |
| – w/DNS anomalies | 2.3K (0.05%) |
| – w/SEQNO anomalies | 9.8K (0.20%) |
| – w/TTL anomalies | 17K (0.35%) |
| – w/RESET anomalies | 8.4K (0.17%) |
| – w/Blockpages | 1.5K (0.03%) |

**Table 1: ICLab dataset characteristics.**

**Block pages.** Finally, the platform analyzes the responses received to identify blockpages that are returned by a censor. This is done by performing regular expression matching with known examples of blockpages (provided by the OONI project [27]) and by comparing responses with those obtained from censor-free vantage points within the United States. In the latter case, we employ techniques developed by Jones *et al.* to identify block pages [19]

**Network paths.** In addition to gathering the above data, the platform also records traceroutes from vantage points to the corresponding destinations of each test.

In total, we utilize 4.9M measurements (with 39K total identified anomalies) from the platform between vantage points located in 539 different ASes (in 219 countries) and 774 URLs. A summary of the ICLab data that we use in our work is summarized in Table 1.

**Ethical considerations and limitations.** To mitigate risk, the vast majority of ICLab's vantage points are obtained via commercial VPN providers (many of which are located in ASes classified as content ASes by CAIDA[5]). This allows us to obtain widespread continuous measurements, without putting users in specific regions at risk. A potential limitation of this decision, is the inability to observe the same filtering as ASes providing residential connections.

In collaboration with the Citizen Lab, we have worked to deploy a handful of Raspberry Pi nodes running the measurement software. Prior to deploying a node, we discuss with the volunteer about the potential risks and they are further presented with a form that summarizes risks for their given country based on existing metrics (*e.g.,* Freedom House [17]). Since the platform does not collect personally identifiable information, our IRB has determined that this project does not constitute human subjects research. Regardless, we maintain contact with any volunteers and monitor the political situations in different regions. Some regions have been deemed too risky to operate in (*e.g.,* Iran, Syria). In general, we aim to balance risk with potential benefits of the measurements.

# 3 LOCALIZING CENSORS

At a high-level, our approach works as follows: First, we use the traceroutes gathered by the ICLab platform to construct boolean clauses such that the literals in the clauses represent ASes observed in the traceroute. We then use the censorship measurements associated with the corresponding traceroutes to assign truth values to the clauses. Finally, the clauses are converted to Conjunctive Normal Form (CNF) and used as input to an off-the-shelf SAT solver. The process is repeated for each type of censorship measurement (*i.e.,* DNS injection, HTTP tampering, and blockpage detection) and various time slices (*i.e.,* for all measurements performed during the same day, week, and month). Next, we analyze the solutions returned by the SAT solver for each CNF. In cases where there are multiple solutions – *i.e.,* multiple truth assignments for a given CNF formulation – we return all literals (ASes) having *True* assignments as potential censors. In cases where there is a single solution, we return all literals (ASes) having *True* assignments as censors. Finally, we characterize censorship leakage by identifying ASes that observe censorship only when they transit through censoring ASes.

## 3.1 Constructing a SATisfiliability problem

Each record in the ICLab dataset contains: (1) the vantage point AS, (2) the URL being tested, (3) the anomaly being tested (and whether it was detected or not), (4) three traceroutes between the vantage point and the URL at the time of testing, and (5) the time at which the test was performed. We use each of these records to create boolean satisfiability problems as follows:

**Clause formulation.** First, we use historical IP-to-AS mapping from CAIDA [6] to convert the IP-level traceroutes to AS-level paths. Next, we eliminate cases with inconclusive paths – *i.e.,* cases where one of the following situations occurred: (1) IP-to-AS mapping was not possible for any of the IPs observed in the traceroute, (2) traceroutes were not possible due to errors, (3) AS-inference was not possible due to non-responsive hops and different ASes observed in the previous and subsequent responsive hop, and (4) there was more than one AS-level path obtained after conversion of the three traceroutes. Each of the remaining AS-level paths forms a clause in our SAT formulation, with each observed AS acting as a literal. The truth value attached to the clause is *True* if the measurement detected its corresponding anomaly, and *False* otherwise. For example, if the AS-level path $X \rightarrow Y \rightarrow Z$ observed DNS censorship, it is represented by the clause $(X_{DNS} \vee Y_{DNS} \vee Z_{DNS}) = T$.

**Time- and URL-based splitting.** Our formulation accounts for the fact that censorship policies and techniques may change over time (*e.g.,* Iran is known to increase censorship during political events such as elections [2]). Not doing so introduces the possibility of generating an unsolvable CNF in the event of a policy change – *e.g.,* the measurement $(X_{DNS} \vee Y_{DNS} \vee Z_{DNS}) = T$ is observed on Day 1 and $(X_{DNS} \vee Y_{DNS} \vee Z_{DNS}) = F$ on Day 2. We address this problem by creating CNFs at four time granularities – days, weeks, months, and year. Additionally, since not all URLs being tested are subject to censorship, we further restrict the CNFs to only include clauses containing measurements to

a single URL. Therefore, we generate one CNF per URL per time granularity (day, week, month, and year). Finally, each CNF is solved using and off-the-shelf SAT solver.

## 3.2 Analyzing SAT solutions

Given a CNF, a SAT solver may return no solution, a single solution, or multiple solutions. When no solution is returned, there is no possible truth assignment to ASes that can satisfy the input CNF. These scenarios may arise due to (1) noise in the ICLab measurements – *i.e.,* incorrect anomaly detection or path inference or (2) changing censorship policies within the specified time granularity. A single solution implies the presence of exactly one satisfying truth assignment to ASes. This ideal scenario allows us to exactly identify ASes that are responsible for generating the measured censorship related anomalies (*i.e.,* the ASes that are assigned a *True* value in the solution). We label these ASes as *censoring ASes*. Finally, when the CNF does not contain enough clauses to generate a single solution, multiple satisfying assignments may be possible. When this situation arises, we consider every AS as a *potential censor* unless the literal associated with it is assigned a *False* value in all returned solutions.

## 3.3 Identifying censorship leakage

In order to prevent leakage of censorship (*i.e.,* where regional censorship policies impact users outside the region), censorship policies need to be implemented in ASes that are either stubs or provide transit services only for ASes within the region. To uncover instances of censorship leakage, we use the following approach: First, we only consider all AS-level paths used in CNFs that return exactly one solution. Next, we identify ASes that (1) are assigned a *False* truth value in the returned solution, (2) are located upstream from the identified censors (*i.e.,* closer to the vantage point being used by ICLab), and (3) are located in a different country from the censoring ASes in the CNF. We label these ASes as victims of censorship leakage due to their inheritance of censorship from censoring ASes in other countries. This inheritance occurs due to their traffic transiting through censoring ASes.

# 4 EXPERIMENTAL RESULTS

We focus on measuring (1) how often our approach is able to generate solvable SAT instances, (2) the amount of path churn observed and its impact on our SAT instance solvability, and (3) the ASes responsible for implementing and leaking censorship.

**Satisfiability of generated CNFs.** As discussed earlier, the CNFs generated by our approach may return (1) no solutions – indicative of changing censorship policies or noise in ICLab measurements, (2) exactly one solution – the ideal scenario, or (3) many solutions – indicative of insufficient number of measurements through diverse paths. In order to understand which scenario occurs most frequently, we analyze the results returned by our SAT solver for CNFs of different time granularities and anomaly types. We find that on average, nearly 92% of our CNFs return exactly one solution and less than 6% of our CNFs return no solution. This indicates high fidelity in our underlying data and highlights our ability to exactly identify censoring ASes.

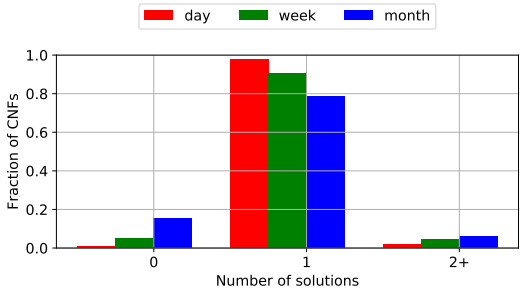


**(a) By CNF granularity.**

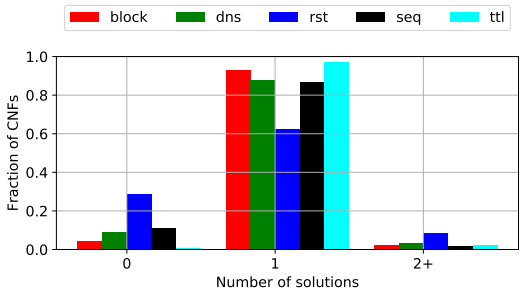


**(b) By anomaly type.**

**Figure 1: Number of solutions found for constructed CNFs when split by CNF granularity and anomaly.**

Our results, when considering CNFs generated for different time granularities and anomalies, are illustrated in Figure 1. Figure 1a shows that as our CNF granularity becomes coarser, its solvability reduces. This is expected since (1) censorship policies are more likely to change and (2) we are more likely to include a noisy measurement in our CNF form when considering larger time periods. Figure 1b shows that nearly 30% of the CNFs generated to identify ASes performing RST injection are unsolvable. This is indicative of low fidelity RST injection measurements from the ICLab platform, due to the difficulty of differentiating between organic and injected RST packets.

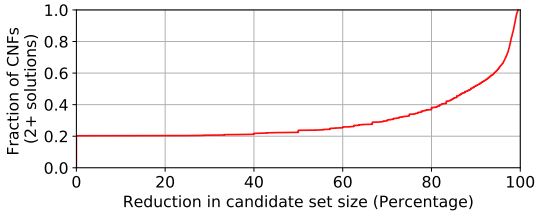


**Figure 2: CDF of reduction in number of potential censors in CNFs with 2+ solutions.**

Even if the CNFs generated by our approach yield more than one solution, they can be useful to identify ASes that

*could not* have been responsible for implementing censorship – *i.e.,* ASes that were assigned a *False* truth value in every solution. We further investigate the 3% of scenarios where constructed CNFs yield more than one solution to identify the impact of this reduction. We find that in 20% of such cases, the solutions satisfying the CNF do not allow for any elimination of ASes as censors – *i.e.,* we are unable to narrow down the set of possible censors by eliminating definite non-censors. However, on average, 95.2% of all ASes in a CNF are identified as definite non-censors, leaving only 4.8% of the observed ASes as potential censoring ASes. Figure 2 shows that 50% of all generated CNFs with multiple solutions have nearly 90% of their ASes eliminated as potential censors.

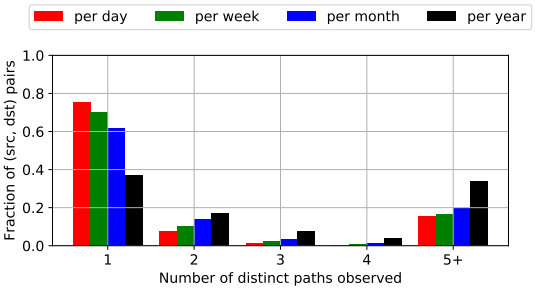


**Figure 3: Number of distinct paths observed between a source and destination AS over varying time periods.**

**Impact of path churn.** Now, we measure the amount of network-level path churn observed over the course of our measurements and consider the impact it has on the usability of our approach. In order to measure the amount of network-level path churn, we measure the number of distinct network-level paths observed between each source (ICLab vantage point) and destination URL for each day, week, month, and year. Figure 3 shows the fraction of these (source, destination) pairs that observe path changes over varying time periods. We find that nearly 25% of all pairs observe path changes within a single day. This fraction increases to 30%, 38%, and 67% when considering periods of one week, month, and year, respectively. Over the period of one year, 35% of the measured pairs recorded at least five distinct network-level paths. Using CAIDA's AS classification database [5], we found no significant differences in the amount of churn observed when considering specific classes of destination ASes (content, enterprise, or transit AS).

To understand the impact of network-level path churn on the effectiveness of our approach, we analyze the solvability of CNFs constructed in the absence of path churn. We eliminate the impact of path churn by only considering the measurements using the first observed distinct path between a source and destination in each CNF. Figure 4 illustrates the number of solutions returned by such CNFs. We see that nearly 80% of all CNFs return five or more solutions (compared to < 1% in the case of CNFs that include multiple distinct paths). We find that although less than 25% of all paths are impacted by path churn each day, the impact

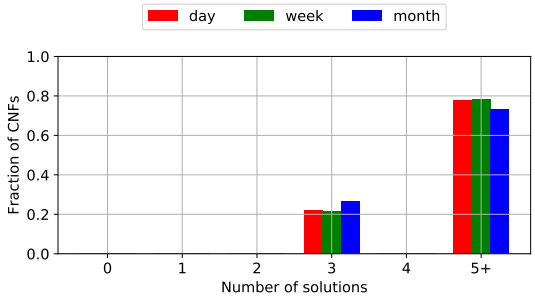


**Figure 4: Number of solutions returned by CNFs in the absence of network-level path churn.**

of these path changes on the solvability of the constructed CNFs is significant.

**Uncovering censors and censorship leakage.** We now analyze the censoring ASes identified by our approach. In total, we identify 65 censoring ASes located in 30 different countries. The countries with the most number of censoring ASes are reported in Table 2. In our analysis we observe a few regions implementing a wide array of censorship approaches. In particular, we find that censors in China and Cyprus implement all measured forms of censorship. Further, the ASes AS1299 (Telianet, Sweden), AS59564 (UNIT-IS, Ukraine), and AS8966 (Etisalat, UAE) are all found to implement at least four of the five measured censorship approaches. Using the McAfee URL categorization database [25], we find that URLs that are most commonly censored fall in the Online Shopping and Classifieds categories. Further analysis reveals that most ASes perform censorship exclusively on few categories of sites, with the exception of ASes in Cyprus which censor content across many different categories. Interestingly, we also identify several ASes (located in Ireland, Spain, and the United Kingdom) which exclusively censor URLs associated with popular ad vendors.

| Region | Censoring ASes | Anomalies |
|---|---|---|
| China | AS4132, AS4812, AS4837, AS17621, AS37963, AS58461 | All |
| United Kingdom | AS5413, AS8928, AS9009, AS20860, AS35017, AS42831 | Block, TTL |
| Singapore | AS4657, AS7473, AS17547, AS38001 | SEQ, TTL |
| Poland | AS20853, AS31621, AS42656 | Block, DNS, SEQ |
| Cyprus | AS8544, AS35432, AS197648 | All |

**Table 2: Regions with most number of censoring ASes.**

To identify which of these 65 censoring ASes leak censorship, we record the regions for which they provide upstream transit (§3.3). We find a total of 32 censoring ASes leak their censorship policies to other ASes. Of these, 24 have censorship leakage extending to other countries. Table 3 lists the ASes responsible for the largest number of leaks outside their region. We find six Chinese ASes in the Top 10, with four other ASes from Poland, Japan, Russia, and the UAE completing the list. Figure 5 illustrates the countries most impacted

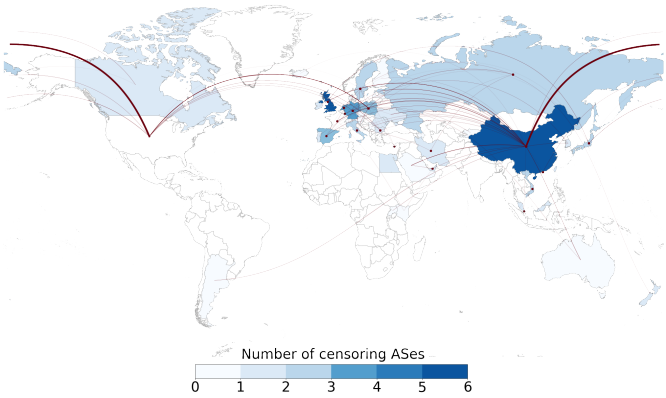


**Figure 5: Flow of censorship. Darker countries contain more censoring ASes and thicker lines indicate more leakage. Sources of leakage are marked with a point.**

by censorship leakage and those containing censoring ASes. We see that with the exception of China, most other leakage is regional – *e.g.,* European and middle-eastern censors leak censorship mostly to other countries in the same region.

| AS | Region | Leaks (AS) | Leaks (Country) |
|---|---|---|---|
| AS58461 Hangzhou-IDC | China | 49 | 21 |
| AS37963 Alibaba-CN | China | 36 | 19 |
| AS31621 QXL-NET | Poland | 28 | 13 |
| AS4812 Chinanet-SH | China | 16 | 9 |
| AS4134 Chinanet-Backbone | China | 12 | 8 |

**Table 3: Censoring ASes with the largest number of censorship leaks in terms of ASes and countries.**

## 5 CONCLUSIONS

In this work, we leveraged boolean network tomography and censorship measurements obtained from the ICLab censorship measurement platform to identify the ASes responsible for inducing censorship related anomalies on the Internet. Our results show that even in the absence of strategically selected monitors and vantage points, exact identification of censoring ASes is possible due to network-level path churn. Our approach uncovered 65 censoring ASes located in (30 different countries) of which 24 were found to leak censorship into other countries. In cases where exact identification of censors was not possible, we were able to reduce the number of potential censoring ASes by over 95%.

The results obtained in this work also uncover the need to improve the fidelity of the RST anomaly detection technique used by ICLab and traceroutes gathered by the platform. In addition to improving the robustness of ICLab measurements (*e.g.,* by using tools such as InTrace [30] in conjunction with standard traceroutes), we also plan to use our approach to extend the ICLab censorship measurement platform in several ways. Specifically, we plan to (1) incorporate data obtained from external performance measurement datasets (*e.g.,* data from M-Lab [23]) to identify ASes responsible for throttling the bandwidth made available to specific protocols used for censorship circumvention and (2) identify, at scale, the ASes responsible for blocking access to Tor bridges [15].

## REFERENCES


[1] Anonymous. 2014. Towards a Comprehensive Picture of the Great Firewall's DNS Censorship. In *4th USENIX Workshop on Free and Open Communications on the Internet (FOCI 14)*. USENIX Association, San Diego, CA.

[2] Simurgh Aryan, Homa Aryan, and J. Alex Halderman. 2013. Internet Censorship in Iran: A First Look. In *Presented as part of the 3rd USENIX Workshop on Free and Open Communications on the Internet*. USENIX, Washington, D.C.

[3] Iljitsch Van Beijnum. 2010. China censorship leaks outside Great Firewall via root server. https://arstechnica.com/tech-policy/2010/03/china-censorship-leaks-outside-great-firewall-via-root-server/. (2010). Online; accessed June 2017.

[4] Martin A Brown. 2008. Renesys blog: Pakistan Hijacks YouTube. http://dyn.com/blog/pakistan-hijacks-youtube-1/. (2008). Online; accessed June 2017.

[5] CAIDA. 2017. AS Classification. http://www.caida.org/data/as-classification/. (2017). Online; accessed June 2017.

[6] CAIDA. 2017. IPv4 Routed /24 AS Links Dataset. http://www.caida.org/data/active/ipv4_routed_topology_aslinks_dataset.xml. (2017). Online; accessed June 2017.

[7] A. Coates, A. O. Hero III, R. Nowak, and Bin Yu. 2002. Internet tomography. *IEEE Signal Processing Magazine* 19, 3 (May 2002), 47–65.

[8] Mark J Coates and Robert David Nowak. 2000. Network loss inference using unicast end-to-end measurement. In *ITC Conference on IP Traffic, Modeling and Management*. 28–1.

[9] Jedidiah R Crandall, Daniel Zinn, Michael Byrd, Earl T Barr, and Rich East. 2007. ConceptDoppler: a weather tracker for internet censorship.. In *Proceedings of the 14th ACM Conference on Computer and Communications Security*. ACM, 352–365.

[10] Jakub Dalek, Bennett Haselton, Helmi Noman, Adam Senft, Masashi Crete-Nishihata, Phillipa Gill, and Ronald J Deibert. 2013. A method for identifying and confirming the use of URL filtering products for censorship. In *Proceedings of the 2013 conference on Internet measurement conference*. ACM, 23–30.

[11] Jakub Dalek, Bennett Haselton, Helmi Noman, Adam Senft, Masashi Crete-Nishihata, Phillipa Gill, and Ronald J. Deibert. 2013. A Method for Identifying and Confirming the Use of URL Filtering Products for Censorship. In *Proceedings of the 2013 Conference on Internet Measurement Conference (IMC '13)*. ACM, New York, NY, USA, 23–30.

[12] Amogh Dhamdhere, Renata Teixeira, Constantine Dovrolis, and Christophe Diot. 2007. Netdiagnoser: Troubleshooting network unreachabilities using end-to-end probes and routing data. In *Proceedings of the 2007 ACM CoNEXT conference*. ACM, 18.

[13] Nick Feamster, David G Andersen, Hari Balakrishnan, and M Frans Kaashoek. 2003. Measuring the effects of Internet path faults on reactive routing. In *ACM SIGMETRICS Performance Evaluation Review*, Vol. 31. ACM, 126–137.

[14] Anja Feldmann, Olaf Maennel, Z Morley Mao, Arthur Berger, and Bruce Maggs. 2004. Locating Internet routing instabilities. *Proceedings of the 2004 conference on SIGCOMM* (2004), 205–218.

[15] David Fifield and Lynn Tsai. 2016. Censors' Delay in Blocking Circumvention Proxies. In *6th USENIX Workshop on Free and Open Communications on the Internet (FOCI 16)*. USENIX Association, Austin, TX.

[16] Arturo Filasto and Jacob Appelbaum. 2012. OONI: Open Observatory of Network Interference.. In *FOCI*. 2nd USENIX Workshop on Free and Open Communications on the Internet.

[17] Freedom House. 2016. Freedom on the Net 2016. https://freedomhouse.org/report/freedom-net/freedom-net-2016. (2016). Online; accessed

June 2017.

[18] Umar Javed, Italo Cunha, David Choffnes, Ethan Katz-Bassett, Thomas Anderson, and Arvind Krishnamurthy. 2013. Poiroot: Investigating the root cause of interdomain path changes. In *Proceedings of the ACM SIGCOMM 2013 conference on SIGCOMM*. ACM, 183–194.

[19] Ben Jones, Tzu-Wen Lee, Nick Feamster, and Phillipa Gill. 2014. Automated detection and fingerprinting of censorship block pages. In *Proceedings of the 2014 Conference on Internet Measurement Conference*. ACM, 299–304.

[20] Ethan Katz-Bassett, Colin Scott, David R Choffnes, Ítalo Cunha, Vytautas Valancius, Nick Feamster, Harsha V Madhyastha, Thomas Anderson, and Arvind Krishnamurthy. 2012. LIFEGUARD: Practical repair of persistent route failures. *ACM SIGCOMM Computer Communication Review* 42, 4 (2012), 395–406.

[21] Sheharbano Khattak, Mobin Javed, Philip D. Anderson, and Vern Paxson. 2013. Towards Illuminating a Censorship Monitor's Model to Facilitate Evasion. In *Presented as part of the 3rd USENIX Workshop on Free and Open Communications on the Internet*. USENIX, Washington, D.C.

[22] Philip Levis. 2012. The collateral damage of internet censorship by dns injection. *ACM SIGCOMM Computer Communication Review* 42, 3 (2012), 21–27.

[23] M-Lab. 2017. M-Lab Data. Overview. https://www.measurementlab.net/data/. (2017). Online; accessed June 2017.

[24] Liang Ma, Ting He, Ananthram Swami, Don Towsley, and Kin K Leung. 2017. Network capability in localizing node failures via end-to-end path measurements. *IEEE/ACM Transactions on Networking* 25, 1 (2017), 434–450.

[25] McAfee. 2017. Customer URL Ticketing System. www.trustedsource.org/en/feedback/url?action=checklist. (2017). Online; accessed June 2017.

[26] Zubair Nabi. 2013. The Anatomy of Web Censorship in Pakistan. In *Free and Open Communications on the Internet*. USENIX. http://censorbib.nymity.ch/pdf/Nabi2013a.pdf

[27] OONI. 2017. Open Observatory of Network Interference. https://ooni.torproject.org/. (2017). Online; accessed June 2017.

[28] Dan Pei, Matt Azuma, Dan Massey, and Lixia Zhang. 2005. BGP-RCN: Improving BGP convergence through root cause notification. *Computer Networks* 48, 2 (2005), 175–194.

[29] Abbas Razaghpanah, Anke Li, Arturo Filastò, Rishab Nithyanand, Vasilis Ververis, Will Scott, and Phillipa Gill. 2016. Exploring the Design Space of Longitudinal Censorship Measurement Platforms. *CoRR* abs/1606.01979 (2016). http://arxiv.org/abs/1606.01979

[30] Robert Swiecki. 2017. Enumeration of IP hops using existing TCP connections. https://github.com/robertswiecki/intrace. (2017). Online; accessed June 2017.

[31] Renata Teixeira and Jennifer Rexford. 2004. A measurement framework for pin-pointing routing changes. In *Proceedings of the ACM SIGCOMM workshop on Network troubleshooting: research, theory and operations practice meet malfunctioning reality*. ACM, 313–318.

[32] Yolanda Tsang, M. Coates, and R.D. Nowak. 2003. Network Delay Tomography. *Trans. Sig. Proc.* 51, 8 (Aug. 2003), 2125–2136.

[33] Y. Vardi. 1996. Network Tomography: Estimating Source-Destination Traffic Intensities from Link Data. *J. Amer. Statist. Assoc.* 91, 433 (1996), 365–377.

[34] Nicholas Weaver, Robin Sommer, and Vern Paxson. 2009. Detecting Forged TCP Reset Packets.. In *16th Network and Distributed System Security Symposium (NDSS2009)*. Internet Society.

[35] Jian Wu, Zhuoqing Morley Mao, Jennifer Rexford, and Jia Wang. 2005. Finding a needle in a haystack: Pinpointing significant BGP routing changes in an IP network. In *Proceedings of the 2nd conference on Symposium on Networked Systems Design & Implementation-Volume 2*. USENIX Association, 1–14.

[36] Xueyang Xu, Z Morley Mao, and J Alex Halderman. 2011. Internet censorship in China: Where does the filtering occur?. In *International Conference on Passive and Active Network Measurement*. Springer, 133–142.