A Quest for Cyber High Ground: Impact of Internet Structure on (Anti-)Censorship

By

Devashish Gosain

Under the supervision of Sambuddho Chakravarty

Computer Science and Engineering

Indraprastha Institute of Information Technology Delhi

New Delhi– 110020

July, 2020

A Quest for Cyber High Ground: Impact of Internet Structure on (Anti-)Censorship

By

Devashish Gosain

A Thesis

submitted in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

Computer Science and Engineering

Indraprastha Institute of Information Technology Delhi

New Delhi– 110020

July, 2020

Certificate

This is to certify that the thesis titled *A Quest for Cyber High Ground: Impact of Internet Structure on (Anti-)Censorship* being submitted by *Devashish Gosain* to the Indraprastha Institute of Information Technology Delhi, for the award of the degree of Doctor of Philosophy, is an original research work carried out by him under my supervision. In my opinion, the thesis has reached the standard fulfilling the requirements of the regulations relating to the degree.

The results contained in this thesis have not been submitted in part or full to any other university or institute for the award of any degree or diploma.

July, 2020

Dr. Sambuddho Chakravarty

Indraprastha Institute of Information Technology Delhi

New Delhi 110020

Abstract

The original design of the Internet was a resilient, distributed architecture, that should be able to route around (and therefore recover from) massive disruption — up to and including nuclear war. However, network routing policies and business decisions cause traffic to be often routed through a relatively small set of Autonomous Systems (ASes). It has practical implications — some of these frequently appearing ASes (*i.e.,* "key" ASes) are hosted in censorious nations. Other than censoring their own citizens' network access, such ASes may inadvertently filter traffic for other foreign customer ASes as well.

Thus in this thesis, we analyzed, **"how can inferences drawn from Internet maps be used to aid (Anti-)censorship?"** Specifically, we attempted to answer questions like — Is Internet structure still hierarchical? What are the key ASes, and in which countries they are located? Can censorious countries (which may host key ASes) filter Internet traffic of other nations transiting them?

To begin with, we constructed a map of the Internet and examined the extent of *routing centralization* on the Internet. We identified the major players who control the "Internet backbone" and point out how many of these are, in fact, under the jurisdiction of censorious countries (specifically Russia, China, and India). We found that approx. one-third of the Internet backbone belongs to the aforementioned known censors that may potentially monitor a large fraction of global Internet traffic. Further, we went ahead to study whether this *hierarchy exists within the nation(s) itself?* If so, can censorious nations exploit this hierarchy to achieve censorship/-surveillance within their national boundary? With censorship mechanisms deployed in a few key ASes, a censor may achieve large scale censorship within its territory. As a case study, we selected India that has the second-largest Internet userbase. We conducted a study on the Internet hierarchy in India from the point of view of the censor. We then consider the question (feasibility) of whether India might potentially follow the Chinese model and institute a single, government-controlled filter. We found that a few "key" ASes (1% of Indian ASes) collectively intercept 95% of paths to the censored sites, and also to all publicly-visible DNS servers. Five thousand routers spanning these key ASes would suffice to carry out IP or DNS filtering for the entire country; 70% of these routers belong to only two private ISPs.

However, the previous feasibility study does not consider the present censorship mechanisms and infrastructure employed by Indian ISPs. Thus, we developed various techniques and heuristics to assess the pervasiveness of censorship and study the underlying mechanisms used by these ISPs to achieve them. We fortified our findings by adjudging the coverage and consistency of censorship infrastructure, broadly in terms of the average number of network paths and requested domains, the infrastructure censors. Our results indicate an apparent disparity among the ISPs — what they filter and on how they install censorship infrastructure. For instance, in

Idea cellular (a popular ISP), we observed the censorious middleboxes in over 90% of our tested intra-AS paths, whereas for others like Vodafone, it is as low as 2.5%. We later devised novel anti-censorship strategies that do not depend on third-party tools (like proxies, Tor, and VPNs, etc.). We managed to access all blocked websites in all ISPs under test.

It must be noted that the proposed anti-censorship solutions were temporary, *i.e.,* they were based on obfuscating the pattern matching used by the censorship infrastructure (*e.g.,* in HTTP GET request changing the Host: evil.com to HoSt: evil.com). If in the future, censors evolve and improve their infrastructure, the proposed solutions may likely fail. Thus, we focused on a relatively new anti-censorship scheme Decoy Routing, which aims to end the arms race between the censor and the free speech activists. Decoy Routing, the use of routers (rather than end hosts) as proxies, is a new direction in anti-censorship research. However, practical decoy routing deployment poses a new challenge of *where to place decoy* Routers (DRs) on the Internet? Thus, we proposed an efficient decoy router placement strategy that requires the construction of global (and country-level) Internet maps. We found that few (≈ 30) ASes intercepted over 90% of paths to the top n sites worldwide, for n = 10, 20...200 and also to other destinations. Our first contribution is to demonstrate with real paths that the number of ASes required for a world-wide DR framework is small (≈ 30). Our second contribution is to consider the details of DR placement — not just in which ASes DRs should be placed to intercept traffic, but exactly where in each AS. We found that even with 30 ASes, we still need a total of about 11,700 DRs.

Decoy Routing requires accessing web content hosted outside the censors' boundary. However, Content Distribution Network (CDNs), which are designed to bring web content closer to end-user, might pose operational challenges to DR. Popular web content (*e.g.,* Alexa popular websites) served from CDNs, might be available within the censor's boundary itself. Thus, we analyzed how do *CDN-based web content localization* can hinder such systems. Moreover, we quantitatively analyzed the impact of CDN localization on various anti-censorship systems, including DR. Such analysis requires geolocating the websites. Thus we adapted a multilateration method, Constraint Based Geolocation (CBG), with novel heuristics and termed it as *Region Specific CBG (R-CBG)*. In $\approx 91\%$ cases, R-CBG correctly classifies hosts as inside (or outside) w.r.t. a nation. Using R-CBG, we observe that most of the popular websites are hosted inside each of the nations. Our empirical study, involving five countries, shows that popular websites ($\approx 80\%$ of Alexa top-1k for each nation) are hosted within a client's domicile. These results reveal that anti-censorship approaches like DR may not directly use a significant fraction of popular websites. However, a small, yet a significant set of websites ($\approx 20\%$), are hosted outside the censors' boundaries and may be used.

Dedication

To,
All those who understand the importance of a question...
And do not consider any question to be a “wrong” question...

Tell people there’s an invisible man in the sky who created the universe, and the vast majority will believe you. Tell them the paint is wet, and they have to touch it to be sure.

George Carlin

In questions of science, the authority of a thousand is not worth the humble reasoning of a single individual.

Galileo Galilei

Acknowledgements

"**At times our own light goes out and is rekindled by a spark from another person. Each of us has cause to think with deep gratitude of those who have lighted the flame within us.**"

Albert Schweitzer

I thank all of those who have helped me in shaping me into what I am — sometimes overtly and sometimes covertly.

Last but not the least, I thank almighty God, if it exists...

Research Papers From This Work

1. Few Throats to Choke: On the Current Structure of the Internet.
   IEEE Conference on Local Computer Networks (**LCN**), 2017.
2. Mending Wall: On the Implementation of Censorship in India.
   EAI International Conference on Security and Privacy in Comm. Networks (**SECURECOMM**), 2017. (Best Student Paper Award.)
3. The Devil's in The Details: Placing Decoy Routers in the Internet.
   Annual Computer Security Applications Conference (**ACSAC**), 2017.
4. Where The Light Gets In: Analyzing Web Censorship Mechanisms in India.
   ACM Internet Measurement Conference (**IMC**), 2018.
5. Too Close for Comfort: Morasses of (Anti-) Censorship in the Era of CDNs.
   Under submission to **PoPETS** 2020. Available on demand.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

Freedom of conscience entails more dangers than authority and despotism.

Michel Foucault

In the twenty-first century, the world relies on the Internet as the most important platform for communication. No wonder freedom of speech over the Internet has been recognized as a human right by the United Nations [2]. On the other hand, not only Governments of China, Russia, Cuba, *etc.* but also notable democracies such as India, South Africa, and Indonesia, have expressed concern about the freedom of speech over the Internet [3]. Some such concerns are of paramount importance for preservation of human rights — *e.g.,* policing child pornography. But such state control of communication channels has been abused to silence the dissent [4], *even in democracies* like India, where freedom of speech is enshrined in the Constitution of India [5]. Thus in this context, it is natural to ask how free and open the Internet is and how robust it is to censorship by these countries.

The DARPA originally mandated that the Internet should be a resilient communication platform that could even survive tremendous damages. Internet routing, relying on connectionless packet switching, helps to route around network disruptions. However, in the twenty-plus years that the Internet has been operated by commercial companies, it has grown tremendously [6]. Today it consists of more than $50,000$ ASes that peer with one another through complex business

relationships (as *customers*, *peers*, or *providers* [7, 8]) to forward each others' traffic[1] to maximize business incentives. As a result, the Internet has a hierarchical structure, *i.e.,* a few key players (ASes) intercepts a substantial fraction of network paths, and potentially, of Internet traffic as confirmed by Subramanian *et al.* [1] in 2001. They reported that the Internet has a backbone and presented a five-level hierarchy of the Internet from the *dense core* (20 ASes) to *stub customer ASes* (8852 ASes). They indicated that the "core" ASes were major ISPs, from the US, Sweden, and France — countries promoting free online communication.

But, these few key ASes, have the advantage of being able to inspect almost all the traffic of the Internet. Thus, a deliberate or inadvertent inspection of users' traffic may lead to serious privacy breaches [9, 10]. Alarmingly, several such key ASes are now headquartered in censorious regimes, like those in Russia and China, that reportedly suppress free speech [11].

Further, Houmansadr *et al.* [8] demonstrated that, in practice, individuals, companies, and even nations have minimal control over Internet routing. Even in the case of China, the world's largest country by the number of Internet users (720 million) [12], forcing to re-route around a tiny fraction of world ASes would lead to massive and costly disruptions.

1. About 44 customer ASes will have to become transit ASes, requiring massive re-organization and investment in their network infrastructure. Presently, China has only 30 transit ASes.
2. Such re-organizations may dramatically increase the load on several Chinese transit ASes. For example, for one particular transit AS, the load may increase by a factor of 2800.
3. The effective latency seen by the Chinese user would increase eight folds.

In other words, the hierarchical structure of the Internet may make large scale re-routing infeasible. Thus maintaining Internet access while avoiding those key ASes of the Internet, which act maliciously, may not be a viable option. It is therefore natural to ask to what extent the Internet routing policies of key players may impact the network access of other ASes, particularly if such key ASes are hosted in *censorious nations*. It has been already demonstrated [13] that DNS queries originated from free nations like the US, when transit China, may get censored by

[1]The existence of such relationships is a constraint on the paths traversed by traffic. For example, if ASes A and B are both providers to AS X, then X will refuse to carry transit traffic from A to B (or B to A).

the Great Firewall of China. In context, this thesis broadly looks into the following research questions:

1. **Where are the key ASes today, and how effective are they at capturing global Internet traffic?**

   The study by Subramanian *et al.* is almost 15 years old [14]. In the meantime, the Internet has grown from 10,000 to 55,000 ASes. How many key ASes are there at present in the Internet? Are any of these likely censorious (or headquartered in censorious countries)? Answers to these questions can be found in Chapter 2.

2. **How much impact do censorious countries have, on network access of other nations?**

   More specifically, are there any key ASes located in censorious countries that also filters the traffic which originate in other nearby ASes (legally beyond the censors' jurisdiction), but which pass through these key ASes? We address these concerns in Chapter 2.

3. **Do key ASes exist within the censorious country itself? How effectively they can implement censorship within the said country (a feasibility study)?**

   Prior studies reveal that in censorious countries like China [15, 16, 17], Iran [18], and Tunisia [19], the Internet has been intentionally organized as a hierarchical structure to aid surveillance and censorship. But, what about Internet structure in countries like India which are ambivalent about Internet censorship? Are there some key ASes that naturally exist due to peering arrangements between ISPs? Can a censor utilize the knowledge of Internet map (of its country), and the key ASes therein, to achieve effective censorship within its boundary? These questions are addressed in Chapter 3.

4. **What are the mechanisms of censorship employed by censorious nation(s)?**

   Internet censorship can be achieved at different network layers *viz.,* network, transport, and application [20]. In democracies like India, is there any evidence of censorship? What type of censorship infrastructure is employed by different ISPs? Is the censorship consistent within and across different ISPs? How can we bypass such censorship schemes? Answer to these questions can be found in Chapters 4 and 5.

5. **What impact do Content Distribution Networks (CDNs) have on (anti-)censorship?**

   CDNs are designed to bring web content closer to end-users. Inadvertently, this may enhance a country's ability to coerce content providers to regulate (or monitor) access within its boundary. On top of that, the obvious solution, *i.e.,* anti-censorship approaches sadly faces a new dilemma. Traditional ones, relying on proxies, are easily discoverable. On the other hand, newer ones (*e.g.,* Decoy Routing) might not work as they require accessing web content hosted outside the censors' boundary. Thus, we quantified what fraction of popular websites are hosted within the censorious nation and might not be directly used by decoy routing. We present our results in Chapter 6.

The gradual growth of the Internet has led to a structure that concentrates substantial routing power in a small number of ASes [7]. Our research validates this "folk wisdom", and demonstrates that it still holds even though the Internet has grown dramatically in the 15 years [1]. There exist about 30 key ASes that have the potential to surveil a significant fraction of Internet traffic. A Similar hierarchical structure can be found within the censorious nations also. For example, in India, 10 ASes cover $\approx 95\%$ of AS-level paths. A censor can exploit this hierarchical structure to achieve a large scale uniform censorship within its geographical boundaries.

However, the proliferation of CDNs paints a different picture of the Internet. CDNs bring web content closer to the end-user (likely in the same country where he/she resides). On one side, this enhances a country's ability to censor (or surveil) web traffic originating within its boundary. Whereas on other, it reduces the routing centralization in the Internet, as a large fraction of web traffic would not cross the respective national boundaries.

Inadvertently, this also poses a new challenge for anti-censorship solutions that rely on accessing web content hosted outside the censor's boundary. Thus through this thesis, we experimentally confirm that, in the future, it would be imperative to build new anti-censorship systems that would seamlessly function even though CDNs have localized the web content within the respective nation's boundaries.

Chapter 2

Few Throats to Choke: On the Current Structure of the Internet

Under observation, we act less free, which means we effectively are less free.

Edward Snowden

The Internet, as originally mandated by DARPA, is a telecommunication network that can survive tremendous damage. As a packet-switching network, it does not require centralized control; catastrophic damage to one part of the network is simply routed around.

However, in practice, the Internet is not a flat network. It consists of a large number (currently about 55,000) of networks, called Autonomous Systems (ASes), which mostly keep their internal structure a black box and enter into relationships (as customers, peers, or providers) with other ASes to forward each others' traffic. The existence of such relationships is a constraint on the paths followed by traffic.

One consequence of such structure, pointed out by Houmansadr *et al.,* [8], is that individuals, companies, and even nations have very limited control over their connectivity to the Internet. Even in the case of China, the world's largest nation by number of Internet users (720 million) [12], and connected to over 850 ASes, choosing to avoid just 2% of world ASes leads to massive

and costly disruption. *E.g.,* 44 ASes in China have to start functioning as transit ASes. (China has only 30 transit ASes, so this is an increase of $\approx 150\%$. The effective latency observed by the Chinese users would increases by a factor of 8.

In conjunction with observing how a small number of randomly-chosen ASes has a surprising amount of power, it must be noted that not all ASes are equal. In the 2001 study by Katz and Rexford [1], the Internet is demonstrated to be a hierarchy of five levels — Dense Core (≈ 20 ASes. Tier-1 providers, nearly a clique); Transit Core (162 ASes. Mostly peer with dense core or each other); Outer Core (675 ASes. Not all closely connected); Small Regional ISPs (950 ASes. Usually have a single provider); Customer ASes (8852 ASes. Stubs — end consumers). Thus, it is natural to ask just how much power the central ASes of the Internet have. In this regard, we look into the following research questions.

- What are the "backbone" ASes of the Internet, and how effective are they at capturing Internet traffic?
  - The study by Rexford *et al.* is fifteen years old; in this time, the Internet has grown from 10,000 ASes to 55,000. How many backbone ASes are there in the current Internet?
  - Are the "backbone" ASes specifically those with no providers (Tier 1), or are other ASes better able to capture traffic?
- How much impact do censorious countries have, on the functioning of the Internet?
  - Are any backbone ASes located in censorious countries? Could they in fact be filtering traffic to other countries?
  - How much collateral damage can censorious countries inflict on "downstream" ASes in other countries (who are technically outside their jurisdiction)?

It must be noted that these questions about the structure of the Internet have important practical implications. Open access to the Internet is an exceptionally powerful resource, and plays a political role in the world; for this reason, free access to information online has been declared a human right by the United Nations [2]. However, there is a tension between free

speech and keeping the commons safe. We suggest that, if in fact the power to monitor or filter all Internet traffic lies in the hands of a few major companies, this may be a cause for concern.

To answer the aforementioned questions, we began by mapping the AS-level Internet paths connecting various ASes to popular websites. Our approach, following Gao *et al.* [21], was to use publicly available BGP routes (obtained from various Internet Exchange Points across the globe [22]), and the relationships between the ASes [7], to construct a directed AS-level graph of the Internet, connecting IP prefixes to all ASes of the world.

We observed that very few ASes — 30 (*viz.*, 0.055% of world ASes) — consistently intercepted over 90% of the world paths to the popular websites we chose, whether we took top-10, top-20, ... or top-200 websites (as per Alexa). *One-third of these "heavy-hitter" ASes were found to have their official headquarters in known censorious nations like China, India and Russia*. This is in stark contrast to the 2001 study where *none* were in censorious countries [1].

A related concern is raised by the observation that Chinese Internet filtering policies inadvertently censor DNS traffic originating outside, but passing through, China [23]. Traffic filtering and monitoring, by the key "heavy-hitter" ASes hosted in a censorious nation, impacts not only "service" traffic from the nation concerned, but "through" traffic from other nations — a clear breach of domain. Our research findings in Subsection 2.3.3 reveal that collateral damage is a serious issue for every one of the nine censorious countries we study, and most of all for China (over 92% of all the network paths that traverse Chinese ASes originate outside the network boundaries of China).

As an aside, we also observed that "heavy-hitter" ASes included not only Tier-1 ASes, but a considerable number of Tier-2 ASes, and that, conversely, several Tier-1 ASes were not heavy hitters; this was a surprise to us, given the general folk belief that the "core of the Internet" consists of Tier-1 ASes (discussed in Subsection 2.4.1).

2.1 Background and Related Work

There are two bodies of work related to our research. The first involves the study of censorship and how it is implemented in various countries around the world. The second is the study of Internet mapping, or more precisely, determining the routes taken by Internet traffic. We discuss both of these areas briefly in this section.

2.1.1 Internet Censorship

Government censorship of the Internet was systematically studied by Zittrain and Edelman [24], in their seminal analysis of filtering by the People's Republic of China. Important early studies were then contributed by Deibert *et al.* [25], Wolfgarten [26], and Dornseif [27], who describe not only censorship policy but also mechanism of filtering as well as anti-censorship measures. Work in the area has since focused on either determining exactly which content is blocked in a given country (*i.e.,* policy) or how such blocking is performed (mechanism).

Several prior efforts emphasize on policy. *E.g.*, authors have explored the censorship in single countries such as China [28], Iran [29], Pakistan [30], *etc.*; Verkamp and Gupta [31] extend this with a survey of censorship across eleven countries. Several projects provide tools to determine censorship policy: ConceptDoppler [32], HerdictWeb [33], CensMon [34], and Encore [35].

Other studies, focusing on mechanism, show a steady increase in the sophistication of both censorship and anti-censorship, from the early work of Clayton [36] *et al.* (TCP reset) and Park and Crandall [37] (HTML response filtering) to the complex arsenal used by China to block Tor, reported by Winter and Lindskog [38]. Our work, in particular, is strongly influenced by two papers in this group: anonymous [23], which raised concerns that *collateral damage* can be caused by the Internet filtering in a nation, and Houmansadr *et al.* [8], who describe the costs of trying to avoid a particular AS.

2.1.2 Internet Mapping

Our work draws heavily on the construction of a map of routes in the Internet. The early work in this area, such as by Govindan and Tangmunrunkit [39], Chang *et al.* [40], and Shavitt and Shir [41], relies on discovering router-level maps using the tool Traceroute, and then uses heuristics to deduce ASes and their connections. However, we make use of the algorithm by Gao [7], which directly extrapolates AS-level paths using public BGP routing data collected by Routeviews [22].

It is reasonable to ask why we did not employ the Traceroute-based approach, *i.e.,* mapping the Internet by sending traceroute probes from various vantage points to IPs in different ASes, when it has been employed by the CAIDA Ark Project [42] and iPlane [43]. Our reason is that, while our approach does have limitations (we are limited by the accuracy and completeness of publicly available routing tables), an approach based on Traceroute is even more limited — by the network locations and availability of the (volunteer) probing nodes; and by the accuracy of IP-to-ASN mapping. We, therefore, chose the Gao algorithm, as being more accurate in computing AS-level paths between any two randomly chosen ASes.

More recently, Claffy *et al.* [44] and Luckie *et al.* [45] have demonstrated improved methods of Internet mapping, which are very accurate in deducing AS relationships (provider-customer, peer-peer). We have therefore taken the relationships they compute, and used this information in finding routes in the Internet with Gao's algorithm.

2.2 Approach and Methods

Our primary question, in this research, is whether a small set of Autonomous Systems actually route all or nearly all of the traffic in the Internet—and if so, to identify these ASes. A high-level overview of our approach is as follows.

1. Collect BGP-level routes in the Internet, to a large set of important targets (such as Google, Facebook, Amazon *etc.*) and construct an AS-level map of the routes connecting all ASes

to these destinations.

2. Identify the "heavy-hitter" ASes on the map — those which appear on a large fraction (nearly all) of the traces.

3. Repeat the experiment with different sets of target sites, to check that the given heavy hitters are general, and not an artifact of the chosen list of target sites.

It is natural to question why we do not directly map the heavy paths of the Internet. Unfortunately, direct information about the magnitudes of traffic flows is not publicly available. As an approximation, we map the paths from all ASes to the most popular websites, which account for much of the traffic in the Internet [46]. This approach does have vulnerabilities — it is quite possible that, for example, we choose the Alexa top-100 websites for our study, and the map we construct is completely different than for the top-200 or some other equally valid set. In order to guard against such a possibility, we perform cross-validation by repeating the experiment with multiple target sets.

It could be questioned, what are Alexa websites and and why we selected them for our analysis. Internet top website lists, such as the Alexa Top N websites, serve the purpose of providing a representative sample of Internet domains in popular use. It is one of the most popular and widely used top list. It is generated based on web activity monitored by the Alexa browser plugin installed by millions of people [47]. Thus, we selected them as destinations for creating maps of Internet, as we believe they are a good representative of popular Internet destinations.

2.2.1 Mapping the Internet

As discussed in the previous section, there are two principal methods of mapping the Internet. The first method, as used by tools such as CAIDA's Archipelago [42], involves the active measurement of the network using traceroute *etc.* Probes are sent along various paths, and the hop-by-hop path is computed, then abstracted to AS-level resolution. The second method is to collect publicly-available routing information, from the BGP announcements of ASes, and to

collate these routes and extrapolate maps of the Internet.

In this research, we have adopted the second method. We build an AS-level Internet map, using the paths connecting popular websites and the various ASes of the Internet, using the approach described by Gao *et. al.* [7]. The approach estimates paths from a given IP or IP-prefix to every AS in the Internet. The inputs to the algorithm are existing BGP RIBs; we use the BGP routing tables collected by the RouteViews project [22] from Internet Exchange Points (IXes), where several ASes peer and advertise their available routes.

Paths directly obtained from RIBs are termed *sure paths*. ASes on sure paths are called *base ASes*. For example, in the (hypothetical) path $192.0.2.0/24 - 1 - 2 - 3 - 4$, each number represents an AS. The path originates at AS4 (right to left) and terminates at AS1, the home AS of the advertised prefix $192.0.2.0/24$. Note that the suffixes of sure paths are themselves also sure paths.

In addition to sure paths, the algorithm computes new ones. This is done by extending sure paths to other ASes to which there are no explicitly-known paths (from the prefix concerned). The extended path must be loop-free, and must satisfy the *Valley-Free Property* [7]. [The AS-level path between two hosts on the Internet is said to follow a "valley free" path, as the path first rises — an AS, then its provider, then a provider of the provider, etc.; peaks — or plateaus, as it crosses through a peering link — and then descends, through provider-to-customer links, until it reaches the destination. There are no "valleys" in the path; no provider-to-customer links between two customer-to-provider links, or vice versa.]

Our original map uses the top-100 most popular websites (as reported by Alexa) as the target WWW destinations; we then perform cross-validation, to check that our results are not an artifact of these sites (as discussed in detail later in this section).

We now explain our Internet mapping process.

- For each prefix, all sure paths (containing all the base ASes) are selected. (These are simply the RIB entries corresponding to the input prefix).

  Next, these sure paths are to be inspected for possible extension to new ASes, provided

they they satisfy the Valley Free property and have no loops.

- The algorithm searches for ASes that share valid business relations with the current end ASes of paths. (Rather than attempt to infer relationships, we directly used the relationships presented by CAIDA [48].)
- One edge is chosen. It is simply assumed that this edge extends the given sure path by one hop.

  Note that we are trying to find a path from an AS to the target prefix, and that extensions of several sure paths might connect the chosen AS to the prefix. Hence there is a need for tie breaking.

  - The algorithm sorts the possible paths, and selects the *shortest* path to the prefix.
  - In case of a tie, the path with minimum *uncertainty* (length of the inferred path extensions) is chosen.
  - If there is still a tie, the path with the higher *frequency index* (the number of times a sure path actually appears in the RIBs) is selected.
- The frequency with which the chosen edge appears in the RIBs, the uncertainty of the extended path, and the new path length, are updated.

An example:

- Obtain a path from RIBs (*i.e.,* a sure path) *e.g.,* $192.0.2.0/24 - 1 - 2 - 3 - 4$.
- Inspect for possible extension to new ASes, provided they satisfy the Valley Free property and have no loops. Assuming the last AS (i.e., AS4) has AS relationships (obtained from CAIDA [49]) with three other new ASes (i.e., AS5, AS6 and AS7). Now all three would be inspected for possible extension. However, assuming *only* AS5 satisfies the valley free and loop free condition, it would be appended to the existing path (by adding some uncertainty to the sure path). The extended path would be $192.0.2.0/24 - 1 - 2 - 3 - 4 - -5$. [It must be noted that, edge $4 - -5$, represents the *uncertain* path that we have added to the existing sure path.]

- Assuming for AS5, we found two paths that terminate at the same prefix 192.0.2.0/24; which path should eventually be selected? The algorithm introduces a few "tie-breakers".
  - Select the *shortest* path to the prefix. *E.g.,* Among the two paths $192.0.2.0/24 - 1 - 2 - 3 - 4$ and $192.0.2.0/24 - 1 - 2 - 3 - 4 - 8$, the former would be selected as it is the shortest.
  - If the two paths have same length, the path with minimum *uncertainty* (length of the inferred path extensions) is chosen. *E.g.,* Among the two paths $192.0.2.0/24 - 1 - 7 - 9 - 10 - 12$ and $192.0.2.0/24 - 1 - 2 - 3 - 4 - -9$, the former would be selected. This is because, with former path the uncertainty value is 0, but with latter path the uncertainty value is is 1 (due to the inferred path extension with AS9).
  - Again, if there is a tie, the path with the higher *frequency index* (the number of times a sure path actually appears in the RIBs) is selected. Among the paths $192.0.2.0/24 - 1 - 2 - 3 - 4$ and $192.0.2.0/24 - 1 - 2 - 3 - 6$, assuming the former one has occurred 15 times (as a sub-path of other AS paths) and the latter one has appeared 10 times; the former would be selected.

2.2.2 Identifying ASes of interest

To select ASes of interest from our map of Internet paths, we take a greedy approach. Ranking the ASes by *path frequency* (*i.e.,* how frequently an AS appears on the paths in the graph), we keep selecting the most-frequent ASes until we achieve a desired level of coverage. We choose 90% coverage as our target, *i.e.,* we select enough ASes to give us a cover of at least 90% of the paths in the graph.

It may be questioned here why we do not follow the standard approach of CAIDA [50], where the "importance" of an AS is determined by its customer-cone size (the total number of its customers, customers of customers, *etc.*). In Section 2.4, we show that in fact customer cone size is poorly correlated to path frequency — the actual metric of our interest — and explain why this is so.

2.2.3 Validation

The most important question regarding our study, is how general its results are. If for example, we find that a small set of “key” ASes dominate routing in the Internet, can we trust this claim, or is it only true for routes to our sample of target sites (Alexa top-100)?

To address this concern, we repeated our experiment for various target sets (Alexa top-10, top-20, top-30 ... top-200 sites) to see if our results remained stable. Finally, we performed direct cross-validation by computing “heavy-hitter" ASes from paths to one set of sites (Alexa top-100) and checking whether they cover over 90% of paths to a different, disjoint set (Alexa ranks 101 to 225).

In this context, we should also consider why we did not simply use our algorithm to plot paths from every AS to every other AS in the world. The reason is that over 85% of the Internet consists of customer ASes [8], which primarily consume content from a small number of providers; the overwhelming majority of computed paths in such an all-to-all map would see almost no traffic. Our map of paths from all ASes to important destinations, in contrast, gives a reasonable picture of the actual paths taken by traffic.

2.3 Experimental Results

In this section, we present our experimental results. First, we consider the map constructed with paths to our original sample, the Alexa top-100 test sites. We then check whether our results remain unchanged as we vary the set of target sites in our test.

2.3.1 Test 1 : Alexa top-100

The most important result we observe is that the frequency with which “heavy-hitter" ASes appear on paths is remarkably top-heavy, not only for our entire sample of test sites (shown in Figure 2.1) as an aggregate, but even for the individual sites tested (shown in Figure 2.2).

The highest-ranked AS, AS3356 (Level 3 Communications), intercepts $1,492,079$ paths

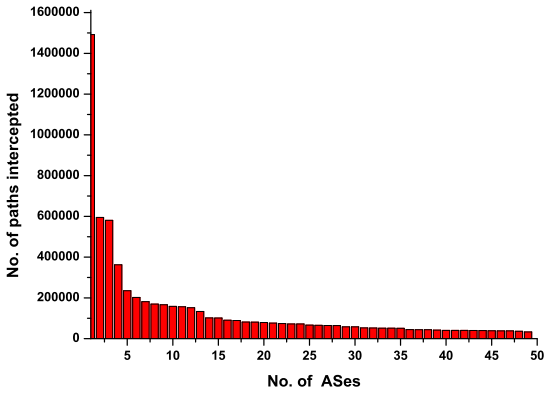

Figure 2.1: Paths to Alexa top-100 sites captured by ASes

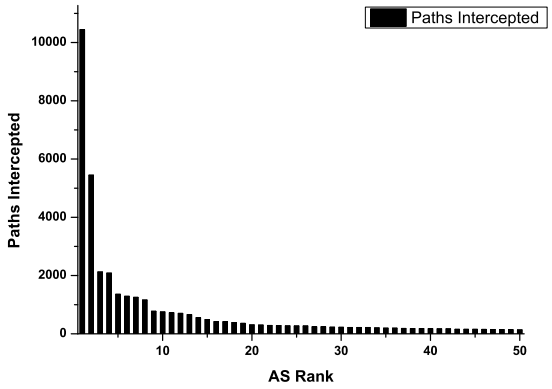

Figure 2.2: Paths to one example target site (facebook.com) captured by ASes

($\approx 33\%$ of total paths).[1] The next highest, AS174 (Cogent Communications), intercepts $536,752$ more paths (not counting overlaps, *i.e.,* paths intercepted by both). Together, AS3356 and AS174 intercept $2,028,831 (= 1,492,079 + 536,752)$ unique IP-prefix-to-AS paths, *i.e.,* about 45% of all the paths. Proceeding similarly, we see that the top 30 ASes by path frequency together intercept 92.4% of all paths. The complete list is presented in Table 2.1. It represents the ASNs, the country where it is headquartered and their ranks with respect to path frequency (P_{freq}) and customer cone size (C_{size}). As is clearly visible, nearly a third of these key ASes lie in censorious countries. (If we include AS 6453, Tata America which, while headquartered in the US, actually belongs to an Indian company — *exactly* one-third of the 30 "key" ASes lies in a censorious country.)

[1] Even this figure underestimates the influence of the company, as another of the 30 key ASes - AS 3549, *i.e.,* Global Crossing—belongs to Level 3.

ASN	Country	Rank (P_{freq})	Rank (C_{size})
3356	US	1	1
174	US	2	2
2914	US	3	5
1299	SE	4	4
3257	DE	5	3
6939	US	6	13
6461	US	7	8
6453	US	8	52
7018	US	9	17
10310	US	10	6
4134*	CN	11	10
3549	US	12	79
4837*	CN	13	85
209	US	14	19
9002	UA	15	97
6762*	IT	16	7
8359*	RU	17	22
2828	US	18	30
20485*	RU	19	21
16509	US	20	9
9498*	IN	21	18
4323	US	22	16
3216*	RU	23	99
2497	JP	24	15
701	US	25	12
12956	ES	26	65
37100	MU	27	23
4826*	AU	28	26
12389*	RU	29	67
1335	US	30	92

Table 2.1: The 30 "key" ASes, which intercept more than 90% of paths. *ASes headquartered in censorious nations highlighted.*

As may be expected, out of the censorious countries, the ones with backbone ASes — Russia (11.09% of world paths) and China (7.39% of world paths), as also India (3.08% of world paths) — cover a substantial fraction of the paths in the Internet[2] (see Table 2.2). This is still much smaller than the U.S. (81.82% of world paths), but overall *censorious nations control* 20.73% *of the paths in the Internet*.

This however portrays a different picture compared to the structure of the Internet that was presented in year 2001 [1]. Approximately 20 ASes that constituted the core of the Internet, were hosted in non-censorious nations (see Table 2.3). Several of these ASes are non-existent now, mostly due to change in business partnerships. Further, newer ASes have gained prominence due to the large fraction of users that they transport. Our results also reflect this massive growth of the Internet — most of the Internet users of the world now reside in Asia; clearly several

[2]In comparison, other censorious nations have much less impact: Iran covers 0.69%, Saudi Arabia 0.23%, and Venezuela, Egypt and Pakistan less than 0.15% each.

"heavy-hitter" ASes are hosted in censorious Asian countries.

Country	Fraction of total paths intercepted
RU	11.09%
CN	7.39%
IN	3.08%
IR	0.69%
SA	0.23%
VE	0.16%
EG	0.12%
PK	0.14%
BH	0.04%

Table 2.2: Fraction of AS level paths intercepted by various countries

Country	ASNs
US	1755, 209, 3356, 4006, 3967, 2914 2828, 7018, 3561, 1239, 8918 6453, 3549, 174, 701, 1, 2548
FR	5511
SE	1833, 4200

Table 2.3: Core ASes of the Internet (as of 2001 [1]) and the countries where they were hosted.

2.3.2 Cross-Validation

In order to verify the generality of our results, we repeated our experiment for various target sets (Alexa top-10, top-20, top-30 ... top-200 sites). In each case we found the same ASes cover $\approx 90\%$ of paths. Further, the key ASes computed using the Alexa top-100, also capture over 90% of paths to the websites ranked 101 to 225 (Figures 2.3 and 2.4). [We add in passing that we also tested how well our "key" ASes covered paths to the 50 most popular non-domestic websites in China, Iran, and Pakistan; they covered $> 90\%$ of these paths as well.]

2.3.3 Collateral Damage

Collateral damage results when an AS filters sites, and also causes its customers to lose access [23]. If, *e.g.*, China was to censor the paths routed through our chosen key ASes, not only would Chinese people would lose access to much of the Internet (and certainly to most popular websites), but also customers of Chinese ASes.

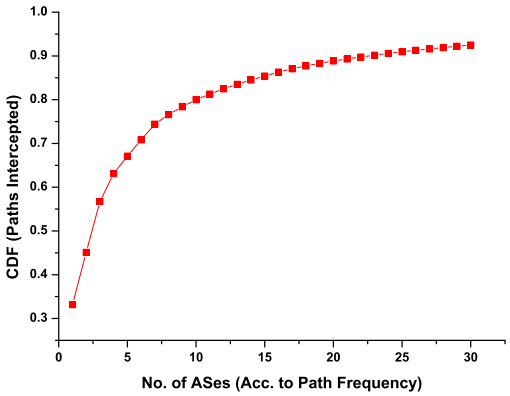

Figure 2.3: Cum. freq.: Paths to Alexa top-100 sites captured by key ASes

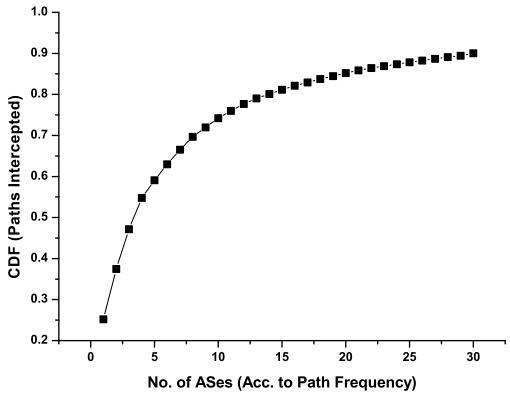

Figure 2.4: Cum. freq.: Paths to Alexa sites 101 - 225 captured by the same ASes

In order to estimate the number of customers that are impacted by such censorship, we inspected the paths through and from nine censorious countries. Figure 2.5 shows the percentage of paths transiting censorious nations that originate at foreign ASes.

We see that in the case of China, *e.g.*, filtering traffic through key ASes would impact many customers, over whom Chinese censorship policies should have no control. $306,874$ AS paths, out of a total of $332,742$ paths involving Chinese ASes and leading to popular destinations, *i.e.*, 92.25% — originate at ASes outside China.

We therefore conclude that the impact of censorship in Russia, China, and India is *not limited to their citizens*. This is especially alarming, considering that China (with $721,434,547$ users) and India (with $462,124,989$) are the largest countries by number of Internet users, and home

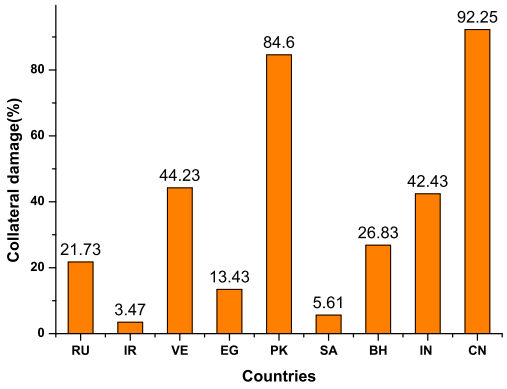

Figure 2.5: Ratio of collateral damage (paths filtered that the country *does not have jurisdiction over*) to intentional damage (paths filtered that actually originate in the country), expressed as a percentage.

to roughly one-third of the world Internet-using population; Russia ($102,258,256$ users) is the sixth-largest, and extremely well connected [12]. In Table 2.4, we show the countries that are potentially impacted by the censorship of China, India and Russia. The list includes most of the major powers: the United States (with $286,942,362$ users), Japan ($115,111,595$), Germany ($71,016,605$), the United Kingdom ($60,273,385$), France ($55,860,330$), *etc.*

Censorious Country	**No. of Countries impacted**
CN	US, RU, IN, JP, PH, NL, GB, LU, DE
IN	US, FR, NL, GB RU, CN, LU, JP, SG, IT, DE
RU	US, IN, CN, JP, NL, CH, GB

Table 2.4: Countries potentially impacted by collateral damage due to filtering by three censorious nations — China, India and Russia

2.4 Discussion and Future Work

From our results in the previous section, it is clear that an overwhelming majority of Internet AS paths in our tests (well over 90%) does in fact pass through one or more of a small set of backbone ASes. This would imply that these ASes have the power to set *de facto* censorship

policy, and monitor or filter Internet traffic worldwide.

The most important question regarding our work, is how we can claim that this picture is true for Internet traffic *in general*, and not an artifact of our methodology, *i.e.,* that the heavy hitters for flows to Alexa top-100 sites are also heavy hitters for flows to *any* site. We have already discussed our answer to this question in the previous sections, with a description of our cross-validation using different sets of target sites. In this section, we also describe our finding that the "heavy-hitter" ASes are not necessarily the Tier-1 ASes.

2.4.1 Heavy-Hitters vs Tier-1 ASes

One of the surprising observations of our research is that the "heavy-hitters" of the Internet not only form a small core, but the size of the core is not much larger than the 20 ASes reported by Subramanian *et al.*, despite the dramatic growth of the Internet in the intervening fifteen years [1].

Another, and perhaps equally surprising, fact is that the "heavy-hitter" ASes we identify are not necessarily Tier-1 ASes (defined as those with only peering relationships, and no providers). For example, our list includes the major Tier-2 ASes Cogent Communications (AS 174) and Hurricane Electric (AS 6939), as well as the ChinaNet backbone (AS 4134 and AS 4837), RosTelecom (AS 12389), Yahoo! (AS 10310) *etc.* which are not only Tier-2 but have Tier-2 providers (Cogent (AS 174) is a provider to RosTelecom, nLayer Communications (AS 4436) to the ChinaNet backbone, and Hurricane Electric (AS 6939) to Yahoo!) We did not, however, observe any Tier-3 ASes. On the other hand, our list does not include five of the sixteen Tier-1 ASes, specifically Deutsche Telekom AG (AS 3320), KPN International (AS 286), Orange (AS 5511), Liberty Global (AS 6830), and Sprint (AS 1239).

We therefore find that the assumption that Tier-1 ASes are the "heavy-hitters" of Internet traffic, is not quite true; there is certainly a strong positive correlation between being Tier-1 and being a "key AS" of the Internet — by which we mean an AS able to intercept most Internet traffic — but it is neither necessary, nor sufficient.

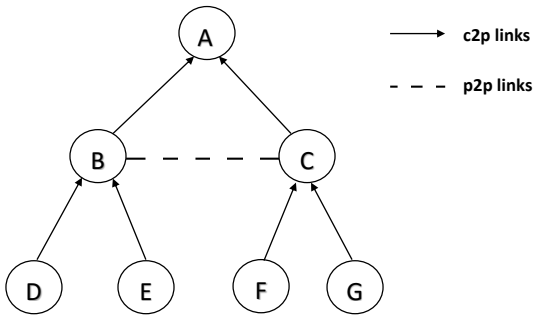

Figure 2.6: Schematic AS graph. A is "root" of customer cone.

Next, we observed that even though many of the ASes on our list were in fact Tier-2, they were very highly ranked by CAIDA [50] in terms of customer cone size. The customer cone for an AS is defined as all the ASes that can be reached via its customers, their customers, and so on.

This raised the question of whether perhaps a composite feature — Tier-1 *or* large customer cone — would predict if an AS is in fact a key AS *w.r.t.* intercepting Internet traffic. However, there are counter examples for this as well, such as RETN (AS 9002) and SOVAM (AS 3216).

We then experimentally checked whether customer cone size is a good predictor of path frequency. Our results were very surprising: in fact, among our key ASes, the Spearman's Rank Correlation Coefficient between cone size and path frequency is only ≈ 0.2. We believe the explanation for this result comes from the existence of *non-root paths* in a customer cone, which we explain with the help of Figure 2.6.

The figure represents a hypothetical AS graph where node A is the "root" AS. A has the highest customer-cone size in this figure (6 ASes - D, B, E, F, C, G). ASes B and C have customer cones of size 2. Many valid (valley-free) paths — such as $D-B-E$, $F-C-G$, $D-B-C-F$, $D-B-C-G$, $E-B-C-F$, $E-B-C-G$ — do *not* pass through the *root AS*, *i.e.,* the node with the highest customer-cone size.

Our map of the Internet shows that this is indeed a common phenomenon. For example, 34.16% of the paths to top-100 IP prefixes traverse the AS with the largest customer cone, AS3356 (cone size = $24,553$). But nearly as many paths, 33.17%, prefer to pass through its 1-hop immediate customers (ref. Table 2.5). For example, as we see in Figure 2.7, the traffic

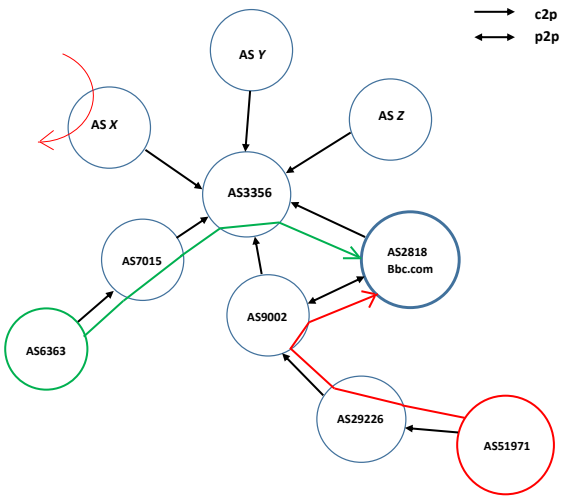

Figure 2.7: Valley free paths in the cone of AS3356. Green line: network path that traverses AS3356 to reach AS2818 directly. Red lines: network paths that traverse the one-hop customers of AS3356, but not AS3356 itself.

through AS9002 to AS2818 (www.bbc.co.uk) does not pass through AS3356, though it is the provider to both these ASes. Still more paths pass through n-hop customers of root ASes (*i.e.*, customers of customers, and so on.) As a result, *customer-cone sizes and AS path frequencies are not well correlated.*

ASN	**% of path not reaching the AS**	**% of path reaching the AS**
3356	34.16	33.17
174	29.05	13.13
2914	28.16	12.90
1299	36.50	8.05
3257	21.00	5.23
6939	7.46	4.40
6461	5.13	4.03
6453	26.00	3.76
7018	7.40	3.70
10310	0.07	3.52

Table 2.5: Fraction of traffic paths in a customer cone traversing large "root" AS, vs fraction traversing 1-hop customers instead.

We conclude that path frequency is not as strongly correlated with customer cone size as we expected, owing to the considerable fraction of paths that do not transit ASes with large cone sizes (preferring to pass through their customer ASes instead). However, for ASes with smaller customer-cones, we observed fewer such non-root paths (possibly because an AS in a small cone

tends to have fewer peers to route through). In future, we may perform a more extensive analysis of such behavior.

2.4.2 AS path estimation

The two main methods of estimating an inter-AS topology are: (a) using `traceroute` traces (as in CAIDA Ark) (b) using BGP routing tables. Traceroute data, being constrained by the location of available probing nodes, is not sufficiently rich to estimate the actual path of traffic from every AS to a given prefix. Hence we choose the routing-table approach.

Previous efforts use simulated BGP paths [51], or paths derived from a Breadth-First traversal of inter-AS links [52]. We improve upon this by employing Gao's algorithm with *real* BGP tables (collected from various Internet Exchanges [22]), thereby estimating the *actual* paths from a chosen IP prefix to all ASes (at a given point of time).

Of course, our map is still not perfect. As Gregori *et al.* [53, 54] and Giotsas [55] show, publicly-available routing tables have biases, errors and bogus route advertisements. Most importantly, they lack the visibility of the "peering" links between the ASes. For example, in Figure 2.6, if the route collector is placed in root AS "A", it would not collect the paths that constitute a peering link B–C (*e.g.,* a path $D - B - C - F$). To address such issues, in future, we may cross-validate our map with different sources of data such as Isolario project [53].

2.4.3 Current and Future Work

The primary idea that motivates this work is to map the Internet, and determine which entities (companies and governments) hold the strategic "high ground" of cyberspace. We explore a complementary research direction in another work:

- The largest nation on the Internet by users, China, is highly censorious. India, the second-largest, is rapidly becoming censorious as well. If in future a Great Firewall of India is built along the same lines as the Great Firewall of China, what might it look like, and what mechanisms might it employ? We study this question in our work [56].

Our results indicate that routing in the Internet is indeed dominated by a few heavy hitters, who therefore enjoy a surprising amount of power. However, several other players in the current Internet economy may also be considered "central" to the Web — the major websites themselves (especially the ones who serve as a platform - most prominently Google and Amazon); root DNS servers; and the major Internet Exchanges (DE-CIX, AMS-IX, LINX, IX.br, DATA-IX and MSK-IX, NL-IX, Equinix, *etc.*) The general question, "who holds the high ground," is thus just as complicated for cyberspace as for the physical world. (The question is very similar to asking: is it the player who controls oil wells who is in a strong strategic position? Or the one with the critical ports on trade routes?) We intend to explore this research direction in detail, in the course of our future work.

2.5 Concluding Remarks

The organic growth of the Internet has led to a structure that concentrates substantial routing power in a small number of ASes. Our research experimentally validates this "folk wisdom", and demonstrates that it still holds true even though the Internet has grown and expanded dramatically in the fifteen years since it was first discovered [1]. However, the *main contribution* of our research is to draw attention to the potential for censorship in this top-heavy structure. A third of the 30 key ASes that form the "heavy-hitter" ASes of the Internet lie in censorious countries (unlike what was presented in 2001 [1]), and they cover over 20% of the Internet paths in our tests. Moreover, from direct examination we see that censorious countries filter (and possibly also monitor) a *substantial* fraction of traffic from other countries.

Further, we also discover that the "key" ASes of the Internet, who carry the overwhelming majority of traffic, are not identical to the Tier-1 ASes (as might be expected from the colloquial use of "Tier-1").

We conclude that while it is certainly understandable that the more powerful routing companies successfully increase their influence over time, perhaps such centralization is effectively making the Internet more fragile as it leads to a small number of "throats to choke". We will pursue this direction further in our future work.

Chapter 3

Mending Wall: On the Implementation of Censorship in India

It's very, very difficult I think for us to have a transparent debate about secret programs approved by a secret court issuing secret court orders based on secret interpretations of the law.

Tom Udall

3.1 Introduction

The current study of Internet censorship is mostly focused on openly censorious countries – China [57, 58, 59], Iran [60], Pakistan [61], etc. Even world-wide studies of censorship [62] essentially focus on countries well known for their censorship. However, in practice, many other countries still implement some form of censorship, which may even be more insidious because citizens are barely aware of it (for example, Sweden [63] and France [64]). In this research, we consider the case of India, a major emerging power with over 450 million Internet users [65] (up from 180 million in 2013, and on track to overtake Europe, which has 520 million users in all). India has been ambivalent about its censorship policy for years [66] (for example, in

August 2015, the government ordered 857 target sites blocked, then backtracked in the face of public outcry [67]), but in context of the fact that *legally*[1] the executive branch in India holds unqualified power to block information, it is natural to be concerned about free speech in India. We begin by asking what policy, and what mechanism the Indian government currently employs; how this might change in future; and what unintended effects such censorship might have on foreign traffic transiting Indian ASes.

Our first step was to formally approach the authorities, by filing a *Right to Information* [68] request (RTI), inquiring about the policies and mechanism the government uses to block content. While the policy itself was confidential, the government was willing to share that the responsibility for filtering lies with individual ISPs, and that they could implement any mechanism they choose[2], as long as they *uniformly comply* with the given censorship policy.

In practice, an ad hoc approach to filtering generally leads to inconsistencies and errors [73], especially during updates [74]. Our initial experiments suggest that this is indeed the case; filtering policies are highly inconsistent across ISPs (see Table 3.1), contrary to the government's expectations as stated in the official response. The current "feudal" approach to policing the Internet in India, *viz.,* allowing ISPs to implement their own censorship mechanisms (which, as we show, do not "strictly adhere" to government diktats), results in inconsistent censorship policy enforcement: *e.g.*, our findings show that users may be able to evade censorship more easily when accessing pornographic sites via Airtel, a large private ISP that screens fewer sites, compared to others such as MTNL.

We next consider the question of how, in future, the government might enforce a unified censorship policy for the whole country. The usual mechanism to enforce a single policy, is to redirect all Internet traffic through a single point of control, where all the traffic can be monitored(this approach has been employed by Iran [60], Venezuela [75], and Saudi Arabia [76]). Even in the case of China, a whole layer of state-controlled ASes must be used to act as a filtering layer that provides Internet connectivity to other ASes [76]. Nearly all the filtering is carried out by two Autonomous Systems - AS 4134 and AS 4812 [17].

[1] Information Technology Act of India 2008 (Section 69A).

[2] IP and URL blacklists [69] are common, but ISPs may choose to employ more invasive techniques, such as DNS Injection Attacks [70] or even IP Prefix Hijacking [71, 72].

Can the government, in future, force all networks to re-route their traffic via a chosen ISP so as to monitor the network? We note that India's Internet infrastructure was grown through a laissez-faire approach (closely correlated with the cellular networking boom), and now consists of ≈ 900 ASes (over 170 of which are ISPs) [77]; it would require a massive effort to redirect all traffic through this new provider. Quite likely, the amount of disruption caused by such a redirection would make it difficult for a democratic nation to implement by fiat.

Might the government implement filtering with the existing infrastructure, without necessarily enforcing traffic redirection? For the existing network, is it possible to find a *small set* of "heavy-hitter" ASes (and network elements in these ASes) that can potentially monitor or censor traffic without too much collateral damage? More formally:

- *Is it feasible to filter/monitor India's Internet traffic? If so, how, and where?* Given that India has over 900 ASes,
  1. Are there a small number of key ASes and routers where the government can intercept most Indian traffic to censored sites?
  2. How does the number of censorious ASes required, vary with the censorship technique — *e.g.,* IP blacklisting, DNS Injection, IP Prefix Hijacking?
- *How much collateral damage will traffic filtering cause?* Internet censorship by an "upstream" AS can lead to inadvertent traffic filtering for its customers. How much impact can Indian censorship have on traffic that simply transits Indian cyberspace?

To answer the above questions we map the AS-level paths from each Indian AS to the potentially censored websites (our test corpus includes not only the sites publicly announced as being blocked, but also others from public resources such as Herdict [78]). We then construct router-level maps within these ASes, using Rocketfuel [79]. Finally, we identify the "key" ASes and routers, *i.e.,* those which appear in an overwhelming majority of paths (and which are, therefore, the logical locations for network filtering).

Our experimental findings reveal that ten ASes cumulatively intercept over 95% of the paths connecting Indian ASes to the sites in our study (i.e., potentially censored sites). Eight of these

key ASes, acting together, can poison $\approx 99\%$ of the network paths leading to DNS resolvers in India (as well as other publicly available services such as GoogleDNS and OpenDNS), thus censoring URL requests. Even more alarming, when we consider another mechanism of censorship - IP Prefix Hijacking - we find five ASes, each of which can individually poison the BGP routes for almost all ASes in the country. Even though the actual number of routers needed for such efforts varies dramatically (from 7 in some ASes, to as high as 1782), overall, a total of less than 5000 routers across all the eight ASes are required for IP or DNS filtering – about 70% of which routers belong to two large private ISPs and any one of five key ASes is enough, if the government resorts to more aggressive measures like IP Prefix Hijack.

Finally, we note that paths that transit Indian ASes but originate outside India form a substantial fraction of the Internet: if India were in fact to adopt a comprehensive censorship scheme in its key ASes, she would censor about 1.15% of *all* Internet paths to the censored sites, worldwide. Thus, the above findings would indicate that, in fact, ordinary Indian citizens should be concerned about censorship, and perhaps start to equip themselves with anti-censorship tools [80]. We begin by discussing the background and related work, in the next section.

3.2 Background and Related Work

The interaction of the Internet with government policy (especially censorship and privacy issues) is a controversial subject [81, 82, 83]. Our case study in this research, India, is a democratic nation, but there is sufficient evidence of Indian censorship [84, 85] that anti-censorship research organizations declare India "partly-free" [86]. For example, the Indian government officially demands that organizations (*e.g.,* Google Inc., Microsoft *etc.*) censor pages deemed objectionable [87].

At present, the government delegates the censorship of traffic to ISPs, as per ambiguous blacklists[3]. This loose approach to censoring traffic leads to inconsistent filtering across ISPs – some users may be able to evade censorship by virtue of their provider ISP. The question arises whether the Indian government *can* impose a centralized filter (as seen in *e.g.,* Iran). Creating a

[3]Several authors have mentioned how these blacklists vary over time [88, 89].

new AS and redirecting through it would have high costs in network disruption, latency, service quality, and so on. But such a process will not be necessary *if* the current structure of Indian Internet is already well suited for monitoring and censorship.

To determine the set of ASes and routers where adversary may install infrastructure for censoring a large fraction of network paths, we generated AS and router-level maps of India. We used such maps to identify such key ASes and routers, and the impact they have.

3.2.1 Background

Our research relies heavily on mapping the structure of the Internet, an area of research called *network cartography* [90]. The Internet consists of routers and hosts, but also has some further structure: the routers and hosts belong to Autonomous Systems, which are independent networks (independent in the sense, they themselves choose who to exchange traffic with). Consequently, Internet mapping proceeds at two levels:

1. *AS-level mapping.* For our research, we required Internet maps representing paths connecting IP address of censored site to various ASes. We thus chose Gao *et al.*'s [21] AS path mapping approach. Their technique uses publicly-available BGP routes (obtained from various Internet Exchange Points across the globe [91])) and the relationships between the ASes [92], and outputs a directed graph of the Internet connecting IP prefixes to all ASes of the world.

   Other AS-level mapping approaches, such as the CAIDA Ark Project [42] and *iPlane* [93], involve `traceroute` probes from various vantage points to IPs in different ASes. Such approaches rely on `traceroute` and are generally limited by the network locations and availability of the volunteered probing nodes; they may not provide the AS-level path between any two randomly chosen ASes.

2. *Router-level mapping.* An AS is not a black box, but contains hosts and routers. Mahajan *et al.* [79] show how the internal structure of an AS can be mapped, by a combination of `traceroute` probes, IP alias resolution[4], and reverse DNS lookups.

[4]Different interfaces of the same router, with different IPs, are called IP aliases.

Powers of the Adversary: Our adversary is a censorious government. The adversary aims to filter Internet traffic, and for this purpose may perform IP filtering, DNS injection/URL Filtering, and IP prefix hijacking attacks. We note that even a government has limitations; for example, it would prefer to implement filtering at a small number of locations, rather than at *every* ISP network in the nation, because of both various political and technical factors (*e.g.,* if changing the blacklist implies wide scale router level re-configuration, there will almost certainly be inconsistencies and failures in enforcement).

3.2.2 Related Research

Much of the study of modern Internet censorship was developed in the context of China [20, 94, 17, 95], particularly the different censorship techniques employed and the network destinations filtered. *E.g.*, Winter *et al.* [95] examine how the Chinese authorities use DPI-capable routers to detect Tor Bridges. Others, such as [96], explored the mechanics of DNS filtering and how China is contributing to collateral damage. A major step forward was made by Verkamp *et al.* [62], who deployed clients in 11 countries (including India) to identify their network censorship activities – IP and URL filtering, keyword filtering and DNS censorship *etc.* Later authors – Nabi *et al.* [61] in Pakistan, and Halderman *et al.* [60] in Iran – demonstrate different methods of censorship employed by their respective regimes, as well as different forms of content blocked. Such studies of censorship in repressive regimes are often limiting, as they require Internet access from almost all network locations inside the country (Nabi *et al.* were able to get access from only five locations, and Halderman from only one).

In this research we take a different direction. While we begin by examining instances of network censorship in our target country (India), our main aim is to determine the *potential* for censorship, in case the regime decides to become more censorious. Specifically, how bottlenecked is the Indian Internet? Is it possible for the adversary to place censors in a relatively small set of ASes and routers, and still filter a large fraction of network paths (and thus potentially users)? - if so, this presents a much lower barrier to entry than monitoring in every AS.

The most relevant related work we are aware of, is Singh *et al.*'s study of how Internet

censorship correlates to network cartography [97]. The authors show a strong correspondence between the *Freedom House Index* [98] of a nation and its Internet topology, and indeed, claim that a nation's network topology is the best indicator of a country's level of freedom. Our work makes use of network topology as well: we use it to determine the "key" network locations (ASes and routers) where the adversary (censorious government) would rationally deploy censorship infrastructure, if its aim was to censor all or almost all Internet traffic in the country, and the impact of such measures on network paths originating both within and outside the nation (but transiting Indian ASes).

3.3 Motivation, Problem Description and Methodology

3.3.1 Preliminary Findings and Motivation

Well-studied censorious countries, such as China, Iran, and Saudi Arabia, tend to have a very clear censorship policy. In contrast, India has a rather ad hoc approach: the government expects all ISPs to (independently) enforce its policies. We find that in practice, traffic filtering is *highly* inconsistent across popular Indian ISPs – the set of blacklisted sites varies by orders of magnitude.

ISP	**Website Categories**							
	Escort (150)	**Music** (100)	**Porn** (50)	**Torrents** (30)	**Social** (20)	**Political** (20)	**Tools** (20)	**Misc.** (150)
Airtel	50, 80, 20	82, 6, 12	1, 49, 0	13, 16, 1	8, 10, 2	2, 15, 3	1, 14, 5	80, 41, 29
Vodafone	24, 87, 39	95, 1, 4	2, 45, 3	16, 11, 3	8, 8, 4	0, 13, 7	4, 11, 5	70, 35, 45
Sify	12, 98, 40	1, 75, 24	1, 48, 1	6, 22, 2	0, 16, 4	0, 15, 5	1, 16, 3	11, 75, 64
NKN	11, 105, 34	57, 33, 10	1, 48, 1	10, 16, 4	4, 12, 4	2, 14, 4	1, 14, 5	65, 56, 29
BSNL	41, 69, 40	68, 12, 20	0, 45, 5	12, 14, 4	7, 10, 3	4, 12, 4	3, 14, 3	88, 27, 35
MTNL	27, 98, 25	81, 2, 17	45, 3, 2	15, 12, 3	9, 8, 3	14, 1, 5	2, 12, 6	73, 23, 54
Siti	23, 99, 28	28, 56, 16	44, 4, 2	14, 13, 3	9, 8, 3	1, 14, 5	1, 12, 7	86, 29, 35
Reliance Jio	0, 123, 27	0, 77, 23	0, 38, 12	2, 26, 2	0, 18, 2	0, 16, 4	0, 15, 5	0, 78, 72

Table 3.1: Censorship trends in India: Some initial results.

To study such inconsistencies, we selected a list of 540 potentially censored websites, divided

into 8 different categories (ranging from escort services, to anti-censorship tools like *Tor* [80]). We then systematically observed the censorship policy in different ISPs, by trying to access our potentially-censored websites through them.

Table 3.1 summarizes our findings. The rows represent the ISPs, columns correspond to the category of site which being filtered, and each entry is a 3-tuple (c_n, o_n, x_n) representing the number of each type of response – *censored*, *open*, and *inaccessible*.

- *Censored:* the ISP intercepted the requests, and responded with an HTML iframe displaying a filtering message (indicating that requested URL had been blocked as per the directions from the Department of Telecommunication).
- *Open:* Websites were accessible without filtering.
- *Inaccessible:* Websites were "down". There was not enough information to determine if the sites were inaccessible due to network or system outages, or requests were deliberately filtered or throttled by the ISP.

For example, we probed 150 escort websites through the Airtel network, and observed 50 to be censored, 80 open, and 20 inaccessible.

We note that the variation of censorship by ISP is quite dramatic: Airtel blocks only 1 out of the 50 pornographic sites probed, whereas MTNL blocks 45.

It is clearly difficult to get hundreds of independent ISPs to correctly comply with censorship orders. The question arises whether, *if* the government decides to enforce a single policy, it is able to do so. So the question arises, *are there a few key bottlenecks in the existing network, where filtering may be carried out*?

3.3.2 Problem Description

In our research we are particularly interested in finding a small set of key locations (ASes and routers) that intercept a large fraction of network paths. More specifically, our questions are as follows.

- Is it possible for the government to monitor/censor a large fraction of Internet traffic by controlling only a small number of network locations (*viz.,* ASes and routers)?
- What fraction of traffic could be filtered, and who would be most affected?
- Would such censorship affect users outside the country as well?

3.3.3 Evaluation Methodology

3.3.3.0.1 Identifying Potential Network locations for IP filtering: In order to estimate the locations for installing IP filtering infrastructure, we built an AS-level map using paths in the Internet, then focused on Indian ASes and their connections. Our map was built using Gao's algorithm [99], which finds AS-level paths to the home AS of chosen IP prefixes (in our case, censored sites) from every other AS in the Internet. The algorithm uses links from known AS paths in BGP routing tables; we obtained tables from a number of vantage points [91].

Unlike other nations, which have an unambiguous list of blocked sites [61], India has no clear censorship policy. We created a corpus of sites blacklisted by various government decrees (as reported by popular media), and also added the sites reported as blocked in India by the crowd-sourced censorship-reporting sites like Herdict [78]. These included social media sites, political sites, sites related to unfriendly nations, and p2p file-sharing sites. Finally, we added to the list the adult sites popular in India (as per Alexa [100]).

We randomly sampled about 100 sites from this corpus. We then computed the paths between all Indian ASes and these prefixes. The ASes appearing in these paths were sorted by frequency of occurrence; we thus selected the few most frequent ones.

Do these ASes appear in paths to other potentially blocked sites as well? To answer such questions, we re-estimated our paths with another set of about 220 sites, chosen from the corpus. The heavy-hitter ASes for this new set of paths were the same as the ones found before.

Intra-AS topology generation: In the second round of experiments, we employed the Rocketfuel

algorithm [79] to compute the router-level paths through 10 heavy-hitter ASes (*i.e.,* major Indian ISPs), then identified the routers which occur in a large fraction of paths (*i.e.,* the heavy-hitter routers in heavy-hitter ASes), as follows.

1. Using `planetlab` nodes, we ran `traceroute` probes to three representative IPs in each prefix advertised by the ASes and by their immediate (1-hop) customer ASes.

   `Traceroute` returned router level paths leading to and out of the said ASes.

2. From the `traceroute` trace, we chose the sub-paths consisting of router IPs advertized by the AS under study (*i.e.,* router within the ASes, identified from [101]).

3. We resolved the aliases (corresponding to the discovered router IPs) with `Midar` [102] alias resolution tool.

4. Finally, from the discovered `traceroute` paths we selected the minimum number of routers which cumulatively intercept a large fraction of the paths. To do this we chose the following heuristic:

   - If total number of edge routers are less than total number of edge and core routers that intercept a large fraction of the paths (over 90%), then we selected the edge routers alone (as the set of edge routers cover 100% of paths through the AS).
   - Else, we selected the "heavy-hitter" (core plus edge routers), appearing in a very large fraction of the paths (over 90%); not all edge routers may appear as often as others (edge and core routers appearing in the discovered paths).

Identifying Potential Sites for DNS based filtering: Another common approach to censorship is to prevent the DNS service from resolving requests. The censor either instructs DNS servers (within its jurisdiction) to filter requests for blacklisted URLs, or installs infrastructure to intercept DNS queries on routers (en-route to DNS servers) and respond with bogus IPs or NXDOMAIN responses – also referred to as *DNS Injection* attack.

Filtering DNS requests, either by simply dropping them, or by responding with bogus responses, could be carried out at the DNS server. However, in a country like India, hosting more

than 55000 DNS servers, distributed across different networks, reconfiguring *all* such servers to filter DNS queries for blacklisted sites would not be easy (besides simple disobedience, there would also be misconfiguration bugs, delays, and network downtime). It would be much more practical to identify a few ASes (and routers therein), that intercept all or almost all the network paths connecting DNS servers to all ASes in the country.

To identify key ASes for DNS injection, we began by identifying the DNS resolvers across all Indian prefixes. We probed IP prefixes of every Indian AS for available DNS servers (UDP port 53) using `nmap`, and noted whether the response was *open*, *filtered*, or *closed*. (*Closed* corresponds to ICMP `'destination port unreachable'` message responses from the destination. *Open* means the client received a meaningful response. *Filtered* indicates that the client received no response[5].)

Each IP, for which we obtained a *filtered* or *open* response, was sent a request to resolve the IP address of some popular WWW destinations (*e.g.,* `https://www.google.com`). Addresses that allowed resolution were added to our list of publicly available DNS resolvers.

Finally, using Gao's algorithm, we constructed a graph of prefix-to-AS paths connecting the IP prefixes corresponding to DNS resolvers, and all the Indian ASes. To find the ASes which would be most effective at DNS injection, we identified ASes at the intersection of a large number of these paths.

Impact of IP Prefix Hijack Based Censorship:

In an IP Prefix Hijacking attack, malicious BGP routers advertise fake AS-level paths[6] in an attempt to poison routes to an IP prefix (see Figure 3.1), thus attracting a large volume of traffic [71, 103, 104, 105, 106].

Prefix hijacking is an extremely aggressive attack, and unlikely to be used in practice; but it has been used in the wild (e.g., blocking of YouTube by Pakistani ISPs [107], and also those involving ConEd (US), TTNet (Turkey), Link Telekom (Russia) among others [72]) and remains viable as an orthogonal way of censoring traffic. So for completeness, we have also considered

[5]This may be due to unavailability or filtering by firewall(s)

[6]Alternatively, router misconfiguration can also lead to similar situations [73].

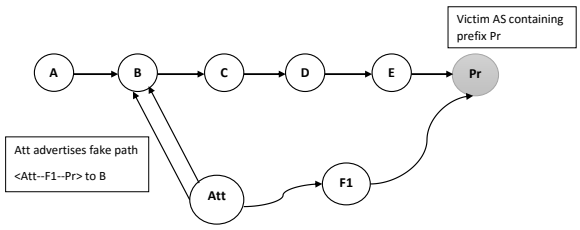

Figure 3.1: IP Prefix Hijacking: Valid path: $A - B - C - D - E - Pr$. A is the origin AS and Pr the AS with the destination prefix. Attacker Att advertises a shorter path $Att - F_1 - Pr$, to AS B. If B chooses this path and directs its traffic to Att, the attacker can censor the traffic.

prefix hijacking as a potential tool for censoring the Internet in India.

In general, for a successful prefix hijack attack, the malicious AS either broadcasts a shorter path to the prefix, or claims to own it outright. The attacking AS advertises fake routes for the targeted prefix to all its neighbors. Ballani *et al.* [71] report that receiving ASes accept these advertisements based on the following heuristics:

1. If there exists a customer path towards the target IP and iff the advertisement presents a shorter customer path, then choose it, else reject it.

2. If there exist a provider path towards the target IP and iff the advertisement presents a shorter provider path, then accept it. For all other cases, the paths are accepted without considering the length.

3. If there exist a peer path towards the target IP and iff the advertisement bears a shorter peer path, accept it. Customer paths are accepted without length considerations while provider paths are ignored.

Estimating the Impact of Prefix Hijack Attack: To study the potential impact IP prefix hijacking, we used the previously constructed AS-level topology and chose an attacker AS with a high *node degree*(*i.e.,* the number of ASes adjacent to the said AS). Inspecting the prefix-to-AS paths, we identified ASes with which the attacker AS had a business relationship, and applied Ballani's heuristics to determine the number of ASes potentially affected by fake advertisements.

Collateral Damage Due to Traffic Censorship: Several non-Indian ASes rely on Indian ASes for Internet connectivity. Censorship activities in Indian ASes may potentially filter the traffic of

these non-Indian customers as well [96]. For example, such unintended filtering was reported by Omantel, that peers with the Indian ISP Bharti Airtel [108]. As one of our research objectives, we try to identify ASes outside India that may be affected by Indian censorship. We identify paths which do not originate in India, but pass through or terminate in India. The non-Indian customers on such paths may face unwanted access restrictions.

3.4 Experimental Results

Continuing from the description of our experiment in the previous section, in this section we present our results. First, we consider router-level filtering, and how many ASes and routers must be selected for effective censorship (in terms of coverage of paths to filtered destinations). Along similar lines, we identify the locations where the adversary could launch a DNS injection attack. We go on to present the results of simulating IP prefix hijack attacks on Indian ASes. Finally, we report the collateral damage to foreign ASes due to IP filtering in India.

3.4.1 Network Locations for IP (Router-Level) Filtering

As mentioned earlier, we first obtained paths connecting Indian ASes to about 100 potential target sites (chosen from our corpus). Figure 3.3 represents the number of paths an individual AS intercepts; the horizontal axis of the graph indicates the ASes, ranked according to the number of paths each one intercepts. A small number of Indian ASes appear in the overwhelming majority of these paths; these ASNs and their owner organizations are presented in the Table 3.2.

The question remains whether the ASes we observe are simply an artifact of the 100 target sites we chose. To check whether this is so, we repeated the experiment with another (non-overlapping) sample of 220 target sites from our corpus. The same 10 ASes covered the vast majority of paths to both sets of target sites, indicating that they are very likely major Indian providers of Internet infrastructure, and cover a majority of paths to *any* target sites.

The cumulative results of *paths intercepted* vs *total number of ASes*, corresponding to both experiments, is presented in Figure 3.2. As evident, *we only need* 4 *ASes to censor over* 90% *of*

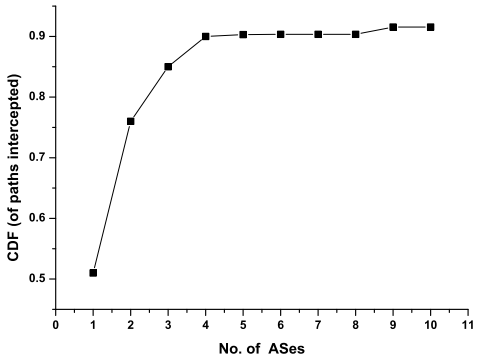

Figure 3.2: CDF of Indian paths intercepted by ASes.

Rank	ASN	Owner
1	9498	Bharti Airtel
2	4755	Tata Comm.
3	55410	Vodafone
4	9583	Sify Ltd.
5	9730	Bharti Telesonic
6	9885	NKN Internet
7	55824	NKN Core
8	45820	Tata Teleservices
9	18101	Reliance Comm.
10	10201	Dishnet Wireless

Table 3.2: AS Ranks, their ASNs and their owners.

the paths to the censored destinations, and 10 *ASes for* 95% *of the paths.* Figure 3.3 represents the number of paths intercepted by each of these ASes individually.

Intra-AS Topology:

We now consider the question of which *routers* (in our key ASes) are responsible for carrying the vast majority of Indian Internet traffic. Following Mahajan *et al.*'s approach [79] (as described previously in Subsection 3.3.3), we create router-level maps of the key ASes, and identify routers that appear on a large fraction of the paths.

Figure 3.4 shows the fraction of paths these routers cumulatively intercept. For privacy concerns, we refrain from revealing the IP addresses of these routers. Table 3.3 represents the number of edge and core routers that cumulatively appear in over 90% of the `traceroute`

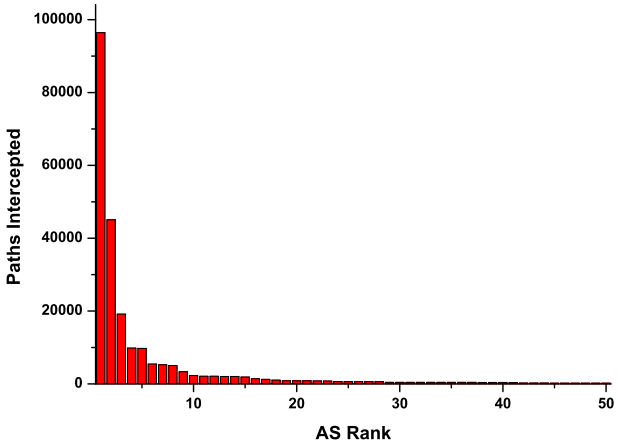

Figure 3.3: Paths intercepted by individual ASes vs AS rank (by path freq.) Total 186679 paths from Indian ASes to 211 prefixes (hosting 320 potentially filtered sites).

paths. The adversary could choose to place filters either at these points - heavy hitter routers of the heavy hitter ASes - or at the edge routers of the ASes, which together see all the traffic that passes through the AS. We find that the total number of edge routers is less than the number of "heavy-hitting" edge and core routers, and conclude that the lowest-cost solution for the adversary is to install censorship infrastructure on the (total of 4996) edge routers.

It must be noted that, at present, the number of key routers varies significantly across ASes, from 7 to 1782. In case of the larger ASes, the AS network administrator could likely improve

ASN	**# of Edge Routers (E)**	**# of Core Routers (C)**	**# of Heavy Hitter Routers (H)**	**# of DR's Required *min*(E, H)**
9498	1782	5321	5192	1782
4755	1779	6229	6434	1779
55410	133	594	634	133
9583	484	4458	4275	484
9730	7	63	62	7
55824	66	325	254	66
45820	193	1147	1132	193
18101	462	2724	2677	462
10201	90	1396	1315	90

Table 3.3: The total number of edge and core routers in 9 ASes that appear in over 90% of the discovered paths. *E.g.,* AS4755 has a total of 8404 routers (1779 edge + 6229 core). However, the total number of edge routers (1779) is less than the number of heavy hitters (6434).

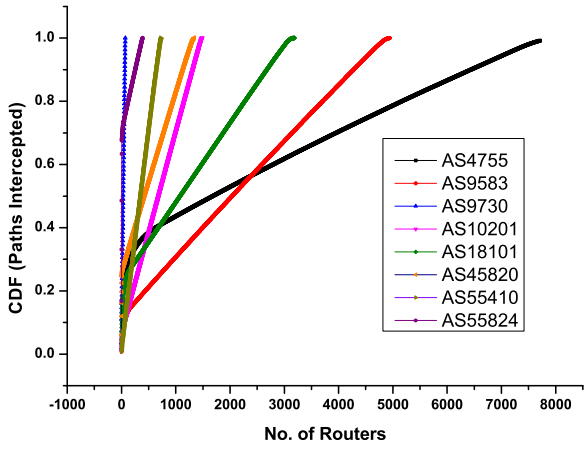

Figure 3.4: CDF of `traceroute` paths intercepted by individual routers, sorted by increasing number of paths through each router (for 8 important ASes.)

on our figures, by combining our findings with better information about the router-level topology and setting routing policy to pass all traffic through a smaller number of routers. Hence our count of 4996 routers is essentially an upper bound, limited by the policies of the present day.

Collateral Damage: Our graph of paths from censored prefixes to ASes has $186,679$ paths of Indian origin (1.76% of paths). A comparable number - $121,931$ paths of foreign origin (1.15% of paths) - transit through or terminate in an Indian AS. *Censorship by Indian ASes may inadvertently impact a very large number of unintended customers, across Finland, Hong Kong, Singapore, Malaysia, the US, and so on.*

3.4.2 Censorship Through DNS Filtering

Using our approach for identifying open DNS resolvers, we identified a total of $55,234$ publicly accessible DNS servers from probing all 12.10 million Indian IPs.

After identifying the prefixes corresponding to these each resolver IP, we selected one corresponding to each AS[7]. In all, we selected 355 prefixes, representative of 355 unique Indian ASes. Finally, using Gao's algorithm, we estimated the paths from each Indian AS to the (prefixes corresponding to) DNS resolvers in India. *Cumulatively, 8 ASes (according to path frequency)*

[7]For multiple prefixes belonging to same AS, we selected one with most resolvers.

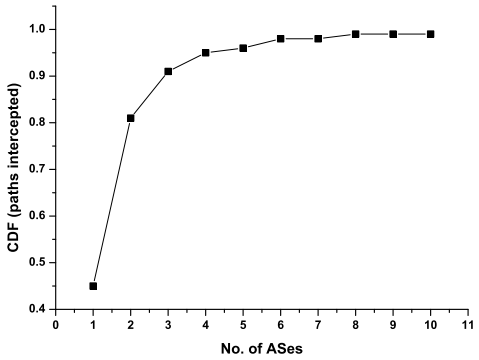

Figure 3.5: CDF of DNS paths intercepted by top 10 Ases.

can intercept 99.14% *of these paths, and potentially launch DNS based filtering or Injection attacks (see Figure 3.5).*

We note that these 8 ASes also appear among the 10 top ASes we identified for IP filtering and IP prefix hijacking. Hence, the same key routers for each of these ASes (as per Table 5.3) may be selected for installing infrastructure to launch DNS injection (or other DNS level filtering schemes). In all, 4906 routers across the 8 ASes can cumulatively filter DNS traffic for all Indian ASes[8].

3.4.3 Censorship Through IP Prefix Hijacking

For IP prefix hijacking, we chose to simulate attacks from the ASes with high node degree. Based on our censored-prefix-to-AS topology graph, we identified the top 10 ASes by node degree, and determined the number of ASes potentially vulnerable to attacks from each of these ASes. The results of these simulations are presented in Table 3.4.

The Table shows that a small number of ASes in India can potentially affect traffic from *all* Indian ASes, as well as a considerable number of foreign ones. For example, fake advertisements by $AS4755$ can impact a total of 955 ASes (896 Indian and 41 others). To effectively launch an IP prefix hijacking attack, the government needs control over the BGP speakers (which form a small fraction of all the routers of an AS); for ASes such as $AS9730$, with 7 edge and 63 core

[8] As mentioned in the previous subsection, this number may be further reduced by routing optimization on the part of the AS network administrator.

		Number of Affected AS'es	
Owner Name	**Attacking ASN**	**Indian**	**Non-Indian**
Bharti Airtel Ltd.	9498	896	59
Tata Comm.	4755	896	41
Reliance Comm. Ltd.	18101	896	41
Vodafone Spacetel Ltd.	55410	896	42
Sify Ltd.	9583	896	58
Bharti Telesonic Ltd.	9730	749	23
Tata Teleservices	45820	560	1
Host Palace	13329	896	45
Dishnet Wireless Ltd.	10201	896	24
Idea Cellular Ltd.	55644	896	37

Table 3.4: IP prefix hijack: A single AS (*e.g.,* AS9498), is well capable of censoring the traffic of all 896 Indian ASes and few (59) non-Indian ASes through prefix hijack attack.

routers, this number is probably very small.

3.4.4 Analysis of Results

We observe that a very small number of ASes (less than 10) intercept a large fraction of AS-level paths connecting Indian ASes to our list of potentially censored sites (obtained from public announcements of censored sites in India), and that this affects a substantial number of foreign users as well. While this result is interesting, there remains the question of whether it applies to censored sites in general, or only the ones in our sample.

Our request to the Indian government, under its own *Right to Information Act* [68], for the complete list of censored sites[9], was refused by the Indian Government Department of Telecommunications and IT, citing confidentiality concerns. Therefore, to cross-validate our results we randomly sampled two sets of target sites from our corpus, and ran our algorithm on each in isolation. The same set of key ASes appeared in both sets[10].

We believe that DNS filtering is a viable threat. Should the aforementioned ASes filter DNS requests, they would also impact over 99% of the AS-level paths connecting Indian ASes to

[9] RTI number: DOTEL/R/2017/50126

[10] We also note that these ASes are, in fact, partners to foreign network providers, and provide connectivity for almost every smaller AS in the country. This is perhaps unsurprising, given the hierarchical nature of the Internet as a whole [92].

DNS resolvers both within and outside India (particularly services such as GoogleDNS and OpenDNS). We note in passing that DNS filtering is more powerful than simple IP filtering: even if a censored site were hosted in a Content Distribution Network (CDN), a user would be unable to reach its content on the CDN, as the request would still have the URL of the origin site, and would thus be filtered.

Finally, while IP prefix hijacking is rarely used (owing to its potential to cause major network outages - *e.g.*, the Pakistan Government's blocking of Youtube [107]), there exist five Indian ASes, each of which could censor traffic for all (or nearly all) Indian users by launching an IP prefix hijack attack. Moreover, only a handful of routers in each of these ASes – *viz.,* the BGP speaking routers may be sufficient for such attacks.

3.5 Limitations and Future Work

3.5.1 Limitations

Our approach in this work is to generate AS and router-level maps of India, and identify the key ASes and routers that intercept a large fraction of network paths. This approach is clearly limited to a snapshot of routing at a moment in time, and in fact we intend to see how our results vary over several years in future work. In addition, our AS-level and router-level mapping algorithms have the following limitations.

AS path estimation (Gao's algorithm): *Our path estimation strategy is limited by the quality of publicly-available BGP routes.*

- *Route-collector bias:* It has been argued by Gregori *et al.* that the existing route collectors (like routeviews [22], BGPmon [109], RIPE [110], PCH [111] *etc.*) miss many of the peering relationships between smaller ASes; our map, as it uses Routeviews data, inherits this weakness.
- *Incorrect route advertisements:* In general, BGP routes are known contain artifacts of misconfiguration and bogus advertisements [112, 107]. Our estimated paths may also be

contaminated with such artifacts.

- *Lack of visibility of peering links:* Many researchers including (Gregori *et al.* [53, 54] and Giotsas *et al.* [55]) highlighted the problem of lack of visibility of peering links on the Internet. The "Isolario" project is one such approach that aims to solve this problem. The goal of Isolario is to improve the knowledge about the AS-level ecosystem of the Internet by increasing the collaboration of different ASes, from which they seek to collect more BGP data. In future, we plan to incorporate their data to improve the AS level topology by adding more peering links

Router level topology estimation: *The discovered topology may not reveal the actual router-level paths for packets traveling between the IPs of the probed AS and the censored websites.*

- *Router-level path variability:* Router-level maps of an AS are far more variable than AS-level maps: the latter rely on AS peering information (which is based on business relationships, that do not change frequently), while the former change with network conditions. Routing tables themselves are prone to inconsistencies and bogus routes [74, 73].

- *Imperfect coverage by Traceroute:* We used a large number of `planetlab` nodes to launch `traceroute` probes [11], but there remains a chance that some routes are simply not covered; further increasing the number of vantage points, *i.e.,* probing hosts, may improve our topology estimation by discovering new paths.

- *Routers filtering* `traceroute` *probes:* In many cases, routers are configured to not reply to `traceroute` probes with the usual `ICMP TTL Expired` messages, and remain anonymous, thereby reducing the accuracy of our estimated router-level topology.

3.5.2 Future Work

Our study of Internet censorship in India can be directly extended to other nations; while our case study was done with Indian data, we make use of no features peculiar to India. We are

[11] The `looking-glass` servers used by the original authors [79] were unavailable at the time of our experiments.

currently extending our analysis to other countries, and developing metrics for how "centralized" a country is (*i.e.,* how many key ASes it takes to censor traffic in a country), as well as how "central" it is in the global Internet (measured by the extent of collateral damage it can cause). There are several other directions to extend this research, which we will explore next.

First - objectionable content is frequently hosted on social media sites, or other sites with apparently benign URLs. Might the government target search engines and social networking sites as well(as seen in China)? Would this be a full blacklist, or partial[12]? And if so, would our key ASes be different for these target websites?

There is also the question of whether popular anti-censorship and anonymity preserving tools like Tor may be attacked by controlling a few network points. Finally, we also intend to consider the question of policing the cellular data network[13], in our future work.

[12] Semantics-based filtering is very hard; *e.g.,* attempts to block jihadi mouthpiece sites also block sites that monitor jihad as a threat, such as `jihadwatch.org`.

[13] As per reports published in recent years, India has 860 million cellular users [113]

Chapter 4

Where The Light Gets In: Analyzing Web Censorship Mechanisms in India

Being a patriot doesn't mean prioritizing service to government above all else. Being a patriot means knowing when to protect your country, knowing when to protect your Constitution, knowing when to protect your countrymen, from the violations of and encroachments of adversaries. And those adversaries don't have to be foreign countries.

Edward Snowden

4.1 Introduction

Free and open communication over the Internet, and its censorship, is a widely debated topic. It is not surprising that an overwhelming majority of prior studies on censorship activities and their mechanism, primarily center around overtly censorious nations like China [57, 58, 59, 17] and Iran [29]. Most of these studies involve reporting censorship activities, with some categorically focusing on the in-depth description of the actual censorship techniques and mechanism that are

employed by such nations *i.e.,* — describing the network location of the censorship infrastructure — what triggers them — and how are clients notified of such filtering.

Through our studies over the past few years, we discovered that even democratic nations like India, have slowly, and rather covertly, evolved an infrastructure for large-scale Internet censorship, involving several privately and federally operated ISPs. India's Internet censorship policies have remained arbitrary (at best ambivalent)[1]. Over time several networks have upped their barriers against users accessing sites, which the administration "believes" to be "unfit for consumption", resulting in enough citizens facing web censorship.

Our previous work [114] emphasized on hypothetical scenarios of potential (future) large scale censorship (or surveillance) by the state. A mere preliminary report was also presented highlighting the inconsistent web censorship policies amongst ASes.

We thus formally approached the authorities, filing a *Right to Information* [68] request (RTI), inquiring about the policies and mechanism the government uses to block content. In response, the authorities shared that while the censorship policies are confidential, the onus of implementing them lied with the individual ASes who could employ any mechanism they chose.

Ambiguous answer from authorities motivated us to conduct our own detailed analysis of the different censorship mechanisms the major network operators of the country employ. We began our research by compiling a corpus of about 1200 potentially blocked sites (PBWs), curated from various Internet sites (*e.g.* Herdict [78], Citizen Lab [115]). Thereafter, we obtained network connections for nine popular ISPs — Airtel, Idea, Vodafone, Reliance Jio, MTNL, BSNL, Siti, Sifi and NKN.

For gauging censorship, we ran the popular censorship assessment tools like OONI [116] on clients hosted in these networks. OONI runs two sets of tests, one at the client and other at their remote control site (assumed to be unfiltered). A mismatch between the results signals potential censorship. However, our initial tests yielded considerably high false positives and negatives when tested for different ISPs. For instance, in Airtel, we obtained a false positive rate of $\approx 80\%$ and a false negative rate of $\approx 11.6\%$.

[1] *E.g.* in August 2015, the government issued orders to block 857 websites, but later backtracked under public outcry [67].

We thus decided to devise our own analysis techniques. We began by observing the ensuing network connection traffic between our client and the censored site. In one particular ISP network, we observed that whenever the client connected to the censored site, it received a valid HTTP response bearing a statutory censorship notification with appropriate sequence number and bits (*e.g.* FIN, RST) in the TCP headers that enforce the client to disconnect with the server. Eventually, the actual response from the censored site also arrives, but by then the connection is already terminated, and the packet is discarded.

All such protocol exchanges hinted toward the presence of malicious network elements (we collectively call *middleboxes*) that snoop (or intercept) users' traffic and upon observing requests to filtered sites, inject the aforementioned crafted packets to censor traffic.

To identify the network location of such censorship infrastructures, we devised a technique which we collectively call *Iterative Network Tracing* (INT), that works on the principle used by `traceroute`. It is quite similar to those proposed earlier by Xu *et al.* [17] and involves sending web requests to censored sites but with increasing IP header TTL values. These messages encounter middleboxes, that are triggered upon the arrival of request to the censored sites.

Using our approach, INT, and various heuristics which we developed after observing peculiarities of censorship techniques, we conducted an investigative study of various censorship mechanism, employed by major ISPs in the country. Our research engages long-term data collections to answer the following questions:

- What sequence of protocol messages triggers censorship?
- Exactly what techniques are employed by ISP networks to filter users' requests to censored sites?
- Approximately what fraction of network paths are impacted by these censors?
- Is censorship uniform and consistent across the various ISPs? More specifically –
  - Do various ISPs block the same set of sites?
  - Do various censorship devices of an ISP (*aka* the *middleboxes*) block the same set of sites?

- How hard or easy is it to bypass such censorship mechanism?

Unlike several previous efforts, that directly draw conclusions based on the results generated by their respective techniques [117, 118, 119] ours, at every possible step, involves corroborating the results via manually connecting to the sites and inspecting the results.

The key contribution of our research efforts, spanning over 18 months, involves detailed answers to all the aforementioned questions. Our findings show that four of these ISPs *viz.* Airtel, Vodafone, Idea and Reliance Jio (potentially carrying a large fraction of network traffic [114]), employ stateful inspection of HTTP requests alone to censors access. For some ISPs, like Idea Cellular, we detected the presence of censorship infrastructure in over a very large fraction (>90%) of the intra-AS network paths. Others, *viz.,* BSNL and MTNL, prime government operators, poison DNS responses for censored sites. In our measurements, we found about 600 censorious DNS resolvers spread across the two ISPs.

Traffic of non-censorious ISPs transiting the censorious ones often gets inadvertently filtered. We observed such phenomenon for various non-censorious ISPs in India. For example, censorship in Vodafone network causes collateral damage to NKN, an otherwise non-censorious educational network.

Through our detailed explorations, we discovered network middleboxes that either intercept traffic (like trans-proxies) or merely snoop on users traffic and sends back specially crafted messages disconnecting the client–server connection. While a vast majority of previous efforts, like [120, 30, 121] report the latter, we discover the presence of intercepting middleboxes (similar to Syria [122]) in one of the ISPs.

Finally, while overtly censorious nations have evolved mechanisms to counter censorship circumvention (proxies and VPNs) [95, 117], we demonstrate simple, yet effective, techniques that can be used to bypass censorship, that do not rely on such proxies, and may go undetected by such censors. Our approach relies on identifying the packets generated from the middleboxes and filtering them at the client, or sending crafted requests that are not detected by the middleboxes, but are correctly recognized by the server. This is harder for ISPs to identify which work by restricting access to anti-censorship infrastructure. Moreover, efforts to retrofit the solution into

existing censorship middleboxes may incur high costs on the part of the middlebox manufacturers and the ISPs, without factoring in downtimes and potential failures.

4.2 Background and Related work

According to ONI report, India is among the list of countries that restricts the Internet content and ranks India as "Partly Free" [123]. Internet censorship in India can be traced back to the year 1999, where website of the popular Pakistani daily newspaper 'Dawn' was blocked from access within India, immediately after the Kargil War [124]. Since then, there are numerous instances of Internet censorship recorded [123] by the orders of the government to an extent of Internet shutdowns. In the year 2015, there were at least 22 instances of Internet shutdowns in different parts of the country [125]. And later in the same year, ISPs were directed by the government to block 857 websites, on the basis of restricting access to pornographic content [126].

Very recently transparency reports published by Facebook [127] and Google [128] confirm that censorship in India is on the rise. It indicates that there were a total of 21 instances of complete Internet shutdowns and 1,228 instances of content removal by Facebook because a majority of content restricted was alleged to violate local laws relating to defamation of religion and hate speech.

Thus, we conducted a detailed study of web censorship trends pertaining to Indian ISPs, specifically aimed to explore the censorship mechanism and its associated infrastructure deployed in the country. We thus begin by discussing important studies in the area of Internet Censorship, primarily reporting the *type* and *mechanism* of censorship. Zittrain [20] in his seminal analysis of censorship observed IP, keyword and DNS filtering in China. Later many studies focused on censorship specific to particular countries *e.g.,* China [28, 57], Pakistan [30], Italy [129], Greece [130] Iran [29], Egypt and Libia [131] etc. Verkamp *et al.* [31] extended this work by deploying clients in 11 countries to identify their network censorship activities encompassing IP and URL filtering, keyword filtering and DNS based censorship *etc.* Gill *et al.* [132] rather than deploying clients, used data gathered by the OpenNet Initiative to detect censorship in 77 countries.

Dalek *et al.* [69] used data from Shodan [133] to identify URL filtering products deployed across many countries including Qatar, Yemen, Saudi Arabia and India. For large scale detection of censorship across multiple countries, there are several projects which provide tools to determine censorship policy: HerdictWeb [33], CensMon [34], Encore [35], OONI [116] and Augur [134].

However, a significant portion of censorship literature focuses only on the People's Republic of China — Great Firewall of China (GFW) [20, 57, 135, 17, 136, 117, 118, 121, 137]. Winter *et al.* [95] studied how DPI-enabled routers detect Tor bridges based on specific TLS cipher suits. Others such as [13] reported that China is heavily contributing towards collateral damage by DNS filtering. Khattak *et al.* [138] observed that GFW operates similarly to NIDS and found exploitable flaws in state management of GFW. Later Wang *et al.* [121] reported that GFW has evolved over a period of time and previous solutions to bypass it [138] are now ineffective. They proposed a novel tool, INTANG, to bypass GFW using carefully crafted packets, without relying on proxies or VPNs.

In the year 2017, we [114] explored that Indian ISPs have incoherent censorship policies and they implement their own content filters resulting in dramatic differences in the censorship experienced by customers. Also, we studied the hypothetical scenario — assuming in future, the Government of India plans to implement strict censorship what would be the probable 'key points' for them to place the filters, for different censorship mechanisms *viz.,* IP filtering, Prefix Hijack, DNS filtering etc.

In this research, we rather carried out a comprehensive study on the '*present*' Internet censorship implementation in India, a vital study missing in our previous analysis. Thus, in this work we developed our own heuristics, with which we tried determining the type of censorship mechanism involved (and in some cases the approximate location of the censorship infrastructure as well). At every step we validated our results against the ground truth, an effort largely ignored by several others in the recent and distant past.

4.3 Data Collection and Approach

For our research, we curated a list of potentially blocked websites (which we refer to as PBW) from different sources which include Citizen Labs [115], Herdict [78] and various past government and court orders of the country [139]. The list includes a total of 1200 websites which we consider to be sensitive (and thus potentially censored). They span across 7 major categories *viz.,* escort services, pornography, music, torrent sites, politics, tools and social networks.

We commenced our research by using the already popular censorship detection tool, OONI [116]. A client which intends to detect possible instances of censorship (at different layers of network stack) installs OONI probe. This fully-automated tool, reports the blocked websites and possible underlying censorship mechanisms being used. After running OONI from five different vantage points we observed that it results in high false positives and negatives. Thus, we created our own scripts to detect Internet censorship in India.

4.3.1 The OONI tool

Open Observatory of Network Interference (OONI) [140], is an open source tool under the Tor project and is designed to detect censorship.

We executed OONI on five popular ISP networks, using the PBW, and recorded the results. To corroborate our findings we also manually checked the sites that were reported by OONI as being censored. To our surprise, we observed very few true positives. An exceedingly large number of sites which were reported as being censored were however easily accessible.

Table 4.1 summarizes our findings. The rows represent the ISPs, columns correspond to the type of censorship reported by OONI, and each entry is a 2-tuple (P,R) representing the precision and recall.

To explain our results better we use the example of data gathered for Airtel (a major ISP of the country[2]). The OONI tool reported that about 78 sites (B_O) were being blocked by the Airtel. Upon manual inspection we observed this number to be much higher, *i.e.,* 133 (B_M).

[2]In terms of network paths it intercepts [114]

Popular ISPs	**Censorship Type**			
	Total	DNS	TCP	HTTP
$MTNL$	$0.57, 0.42$	$0.44, 0.10$	$0, 0$	$0.60, 0.64$
$Airtel$	$0.19, 0.11$	$0, 0$	$0, 0$	$0.19, 0.11$
$Idea$	$0.57, 0.62$	$0, 0$	$0, 0$	$0.57, 0.62$
$Vodafone$	$0.69, 0.82$	$0, 0$	$0, 0$	$0.70, 0.78$
Jio	$0.34, 0.15$	$0, 0$	$0, 0$	$0.36, 0.14$

Table 4.1: Accuracy of OONI: Precision and Recall values, measured in various ISPs.

Only 15 websites ($B_O \cap B_M$) that were actually being censored were also correctly detected by OONI. This provides us with a precision of 0.19 ($|B_O \cap B_M|/|B_O|$) and a recall of 0.11 ($|B_O \cap B_M|/|B_M|$). Similar results were observed for other ISPs as well.

Such low values of precision and recall can be attributed to the fact that OONI tool uses rudimentary approaches to detect potential censoring activities. For instance, while detecting DNS filtering, it compares the IP address of a given host name returned by Google DNS resolver (which they assume to not be tampered) with the IP address mapped to that website by the client's ISP. If the two IP addresses of the same website are different they assume it to be censorship. But, in many cases differences in URL resolution is likely an artifact of network hosting architectures (*e.g.* CDNs).

Also, while detecting HTTP filtering, OONI sends an HTTP request to a given website over the network where the client (running the OONI probe) is hosted. Following that, the same request is sent from the control server (of OONI). HTTP responses obtained from these requests are compared (based on a threshold) and the website is assumed to be filtered if the responses differ. However, while conducting our experiments, we observed that in spite of observing difference in HTTP responses, that OONI uses to report censorship, the websites were actually not blocked. We explain the reasons for incorrect reporting in detail in Section 4.7.

Thus, for our research, we abandoned OONI and created our own *semi-automated* scripts to record the censorship instances by various ISPs across India. For instance, similar to OONI, we tried sending `GET` requests to PBWs from the client in test ISPs and through Tor. If the difference in responses was less than a certain threshold we considered them *non-censored*, otherwise we manually inspect the responses further, unlike OONI, which directly flags them as censored. For

instance, in Airtel network, we observed that for 390 PBWs, the difference between the contents of HTTP responses was more than 30%. On manually verifying (*i.e.,* identifying the censorship message in HTTP response) we confirmed that 156 of those (*i.e.,* 40%) were actually blocked. We repeated the same experiments for several other ISPs and found that $30-40\%$ of the websites which would have been flagged as censored by OONI were actually false positive. Thus, we selected 0.3 as the threshold for our experiments (explained in detail in Subsection 4.3.4).

We now present our approach for determining the type of censorship (DNS, TCP/IP and HTTP) and mechanisms behind them.

4.3.2 DNS Blocking

I. Background: For ordinary netizens, DNS resolution is the primary step for accessing any website. URL entered by such netizens is first resolved to its associated correct IP address. Thus, invariably censors exploit this step, and often return an incorrect IP address, resulting in website's unavailability.

DNS based blocking can be achieved by (1) *DNS poisoning* [141]— whereby a corrupt(-ed) resolver replies with the incorrect IP for specific DNS queries. (2) *DNS injection* [96] — where some middlebox between the client and resolver intercepts the DNS query and deliberately responds with a forged response bearing an incorrect IP address.

II. Identifying sites filtered using DNS requests: In order to identify DNS filtering by ISPs, we selected PBWs that could otherwise be successfully resolved via Tor circuits (ending in exit nodes in non-censorious nations). Thereafter, we attempted to resolve these PBWs through ISPs under test.

A URL might resolve to multiple IPs everytime a resolution is attempted. Further, these responses may also differ when resolution is attempted via Tor and from the ISPs under test (due to reasons such as CDN based hosting). Thus, URL resolutions, attempted via Tor and the tested ISPs, resulting in an overlapping set of IPs, were considered to be uncensored and eliminated from further inspection.

In order to ascertain DNS filtering on the remaining URLs (among the PBWs), we first tested the following intuition: an ISP may deliberately resolve multiple blocked websites to some unique IP addresses. We went ahead and performed DNS resolution for all these remaining URL from the ISP under test.

Invariably, we found several blocked websites resolving to the same IP address belonging to an ISP under test. In general, several sites may resolve to the same IP address, an artifact of modern commercial hosting. Thus, before conducting our measurements we eliminated sites that were actually hosted with the same IP.

Additionally, in several cases the blocked URLs also resolved to bogon IP [142] addresses.

Hence, we applied the following heuristic to decide if the DNS responses were seemingly manipulated by the censor.

1. *Resolved IPs belong to the same AS that hosts the client:* None of the PBWs were hosted in the clients' AS. Thus if any of the resolved IPs belong to the clients' AS (the one under test), the AS is considered censorious and the corresponding URL is marked censored.
2. *Resolved IPs are Bogons*: If any of the resolved IPs is a bogon [142], we consider the AS to be censorious and the URL as being censored.

We applied these heuristics and identified the manipulated IP address responses and removed the corresponding URLs from further analysis. In order to confirm that *only* the aforementioned strategies, and none others, are employed by the tested ISPs, we sent HTTP request to the remaining IPs (obtained initially when resolution was attempted from the tested ISPs), via Tor and manually inspected the corresponding response to confirm that they were as expected.

III. Identifying DNS filtering mechanism: After identifying the set of websites which are censored due to DNS filtering, we intended to identify the mechanism behind the blocking. To that end, we began by identifying all open DNS resolvers of the ISP under consideration. To do this we sent DNS queries requesting resolution for otherwise uncensored sites (*e.g.* our own institution's website) whose correct IP address is known beforehand, to the entire IPv4 address

space of the said ISP. DNS resolvers, even if censorious but otherwise configured correctly, are expected to respond to such queries with legitimate IP addresses.

In order to identify only the *censorious resolvers*, we sent 1200 DNS queries, corresponding to the individual PBWs, to each of the DNS resolvers (in each of the individual censorious ISPs). Resolvers which responded with manipulated IP address (for any of the DNS queries), are considered to be censorious.

To determine **how and where censorship happens**, *i.e.,* if censorious DNS responses were due to middleboxes or by the poisoned resolvers, we used a variant of INT (as shown in Figure 4.1). We began by identifying the router-level path between the client and the censorious resolvers using `traceroute`. Thereafter, the client sends DNS requests (corresponding to PBWs) to *only* censorious DNS resolver by iteratively increasing IP TTL values. Identifying censorship mechanism involved checking if the responses (between the client and a PBW) arrive from any network hop other than the last one. Such responses are likely due to middleboxes, else they are due to poisoned resolvers.

In all our tests we received manipulated IP addresses from the last hop only, indicating the presence of **DNS poisoning**.

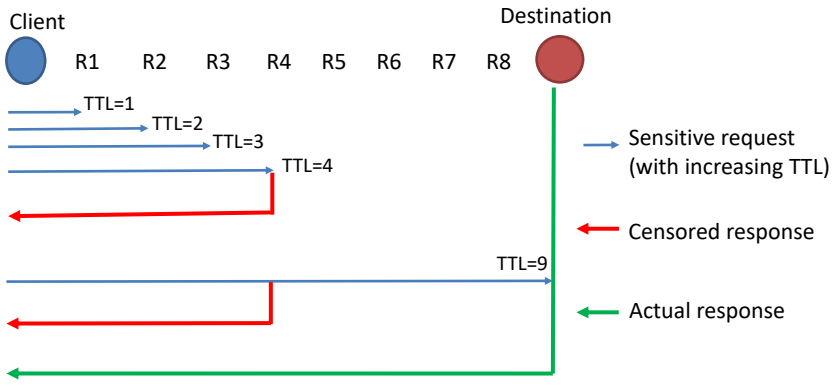

Figure 4.1: Iterative Network Tracer: Client sends a crafted query (DNS query/HTTP `GET` request) containing a blocked domain with increasing TTL. Censorship response is observed from the censor's network element.

4.3.3 TCP/IP Packet Filtering

Network protocol header based filtering is a rather ill defined, albeit very commonly assumed, network censorship technique. It is frequently believed by netizens that ISPs filter traffic based on IP addresses and port numbers. Not surprisingly several past research efforts focused on detecting censorship where authors claimed that ISPs filter traffic based on IP addresses. To that end, they primarily relied on packet drops corresponding to TCP connection attempts [134] and claimed them to be due to IP-level censorship.

In general, such techniques may often mis-classify various kinds of systemic failures such as network congestion, outages and delays in route re-computations, as IP address based censorship. Also, unlike HTTP censorship, which often involves users receiving censorship notification packets, IP address filtering may reveal no information to the client, making it hard to distinguish from the other reasons mentioned. Such scenarios are very difficult to validate (an important, and often ignored aspect of prior research [134]).

Nevertheless, we used a straightforward approach to detect filtering based on network and transport protocol headers. We attempted TCP $3-way$ handshakes with the PBWs. These were tunnelled through Tor circuits terminating in non-censorious countries. For those websites where the connection succeeded (via Tor), we again attempted five subsequent TCP $3-way$ handshakes (from the ISP under test) with a delay of approximately two seconds between each of them. If it failed in *all* the attempts, it implied TCP/IP filtering. *However, in none of the tested ISPs, we ever obtained this form of censorship.*

4.3.4 HTTP Filtering

I. Background: HTTP filtering aims at hampering the communication between client and server by observing content of HTTP packets. The censor can achieve this type of filtering by deploying middleboxes in the network (placed between client and blocked domain).

II. Detecting HTTP filtering: In our experiments, we tested all our chosen ISPs for potential censorship using our curated list of 1200 PBWs. We began by creating Tor circuits terminating in non-censorious countries. Through these, we tried accessing all the PBWs. The retrieved contents were compared against the contents obtained when directly accessing the respective PBWs, from our clients hosted in the individual ISPs.

One may expect that in case of overt censorship the difference between the aforementioned responses to be very high (*e.g.*, >80%). However, there might be false negatives in case the differences are lower. We thus chose a relatively lower difference threshold of about 30%. Manual inspection of site content, when the differences were lower than this threshold, revealed no censorship. Whereas, when the difference was greater than 30%, we found several instances of censorship notifications. In general we observe that whatever threshold we take for measuring the difference between HTTP responses, for the cases where difference is more than the threshold, we need manual inspection. Depending upon the censorship response and ISP under test, the threshold may be adjusted accordingly.

III. Which HTTP messages trigger censorship: We began our study of determining what triggers censorship by observing the protocol messages between client and PBWs. *E.g.*, for a client hosted in Airtel, we observed that after sending an HTTP `GET` request to a PBW (following a regular TCP 3-way handshake), a HTTP 200 OK response packet arrived, whose source IP address is that of the PBW. It had the TCP FIN bit set and payload carried the censorship notification. TCP FIN bit forced the client's browser to initiate TCP 4-way connection termination with the PBW. Eventually, the packet from the PBW *also* arrived. We were thus unsure as to what triggered censorship – requests from the client to the PBW or their responses?

In the past researchers have reported middleboxes in China that inspect both *request* from the client and *response* from the server for censoring content [36, 57]. Thus, we also required heuristics to determine if censorship was triggered by requests or by responses.

In order to distinguish between the two possibilities, we adopted the following approach. Initially, the client runs `traceroute` to obtain the number of hops (n) to the actual website.

Thereafter, the client establishes a TCP connection with the website and sends two consecutive HTTP `GET` requests for the blocked website. The IP header of the first request has a TTL value of $n-1$ and is not expected to reach the site, and thus no responses from the site are expected. Whereas, the second is sent to the site bearing a TTL value of n (and is expected to be handled like a regular request).

Depending upon what the middleboxes en route inspect, there could be three possibilities:

- Possibility 1 (Middlebox get triggered *only* by the request): Both the above requests would traverse the middlebox that would respond back with a censorship notification-cum-disconnection message.
- Possibility 2 (Middlebox get triggered *only* by the response): The middlebox would be triggered when it inspects the response messages, which happens only when the request actually reaches the site and elicits it (*i.e.,* only corresponding to the second request).
- Possibility 3 (Middlebox triggered by request *and/or* response): The middlebox sends censorship notification for both the requests.

In our measurements, we observed censorship notification-cum-disconnection packet for both the requests (*i.e.,* for TTL=$n-1$ and TTL=n). This directly rules out the possibility 2, *i.e* when middleboxes are triggered *only* through responses.

For both the remaining possibilities the middleboxes could inspect both requests and responses. In order to distinguish these possibilities, we crafted our own HTTP `GET` request such that PBW interprets it correctly but *not* the middlebox. *E.g.,* in Airtel network, merely manipulating the case of the HTTP header field `Host` and changing it to `HOST` was sufficient for the request to go undetected by the middlebox (while correctly interpreted by the PBW). We show in Section 4.6 that for all ISPs, we managed to bypass the censorship by only modifying the HTTP header fields of the `GET` request. *This confirms that middleboxes are only inspecting the requests (possibility 1) and not the responses, as otherwise, we would still be receiving censorship notifications when the responses, carrying censored content, would encounter the said middleboxes*.

IV. How `GET` request triggers the middlebox: Since middleboxes inspect the `GET` request for potential censorship, we intend to confirm exactly how the middleboxes get triggered. By default a regular `GET` request bares only the domain name along with the requested page. We first ran `traceroute` to obtain the number of hops to the server. Then, we crafted a `GET` request whose IP TTL was set to the value of penultimate hop, such that it passes the middlebox but never reaches the server. Thus, we ensured that response (if) received is from the middlebox and not the actual server.

In the payload, we fudged the domain name and its offset within the request, to determine exactly what triggered censorship. *E.g.,* we set the HTTP `Host` field to that of an uncensored site, while the domain name of the censored site was positioned at a random offset within the HTTP header (say beyond the requested page indicated in the `GET` field). In all our tests we observed only when the `Host` field is set to the domain name of the censored site, we observed the censorship response. Further details of more related experiments are presented in Section 4.5.2.

As described ahead in Section 4.5, in three of the four ISP where we observed HTTP censorship, the middlebox responds back to the client with a variant of the aforementioned censorship notification-cum-disconnection messages. These were mostly HTTP 200 OK responses carrying the statutory censorship notification along with appropriate TCP bits enabled that force the client to terminate the connection with the PBW. They bore appropriate sequence and acknowledgement numbers (along with other protocol header information) to make them indistinguishable from legitimate packets which the client's underlying protocol stack expects, *w.r.t.* the initial TCP connection to the PBW.

In Section 4.6, we show how we exploit this knowledge of protocol header idiosyncrasies, along with deliberate fudging of the requested domain name in the `Host` field of the `GET` requests to sidestep censorship.

V. Identifying location of HTTP middleboxes: After characterizing the blocking behavior, we intended to identify the *network location* of middleboxes *viz.,* IP address. As earlier in Subsection 4.3.2, we first ran `traceroute` to determine the number of hops between the client and a

PBW (to be tested). We then use INT (shown in Figure 4.1) whereby following a regular TCP handshake to the PBW, we sent series of crafted HTTP `GET` request to it, with increasing TTL values until it encounters the middlebox. The middlebox, upon observing the `GET` request, bearing the domain name of the PBW, responds with a censorship-notification-cum-disconnection message (with TCP FIN/RST bit set). Correlating the TTL value of the request (observed by the middlebox) against the IP address hops reported by `traceroute` helped us identify the middleboxes.

4.4 Experimental Setup and Ethical Considerations

All our experiments were carefully designed to avoid any unintentional or unethical network disruptions or system downtimes and failures, of third-parties, including individuals, ISPs, institutions and governments. To perform our large-scale censorship studies, we used our own client machines hosted in various networks. For this we bought connections of about nine popular Indian ISPs. To augment our results we used hosts deployed outside India – about 50 planetlab hosts (which by default provide "sudo"-able administrative access), about ten virtual machines in various cloud hosting services and around ten more hosts placed in institutions where we had collaborators, who were kind enough to lend us their infrastructure (granting administrative access and all essential privileges to conduct our experiments).

In our initial studies, we sent traffic to a small sample of PBWs, from each of the clients under our control. Using `pcap` we passively inspected the protocol responses to determine the actual mechanism through which censorship was enforced (as explained ahead).

Further, we also conducted large-scale studies to quantify the impact of censorship. Our efforts involved sending HTTP GET requests ($\approx$ 440 Bytes) to Alexa top-1000 sites, from the clients we controlled, at every 8–10 second intervals. This was slow enough to not impact network performance of other users.

4.5 Experimental Results

In order to determine the censorship mechanism we inspected for DNS, TCP/IP and HTTP blocking for the list of PBWs in nine major ISPs of the country (as already explained in Section 4.3). For all ISPs, we found instances of DNS and HTTP filtering only.

4.5.1 DNS Filtering

We began our study by identifying the open DNS resolvers in a chosen ISP. Thereafter applying our heuristics presented in Subsection 4.3.2 we determined which of the 1200 curated PBWs were being censored along with their corresponding DNS resolvers. We observed poisoned DNS resolvers in only two of the nine ISPs, *viz.,* MTNL and BSNL. Before presenting the results, we propose two metrics to analyze the extent of DNS filtering within the ISPs:

1. ***Coverage***: Ideally all DNS resolvers of an ISP must be poisoned. We define *coverage* as the fraction of all the resolvers of the ISP which are poisoned.

2. ***Consistency***: Ideally the same set of sites must be blocked by all the *poisoned* resolvers of an ISP. We determined the set of filtered URLs as well, as all the resolvers that blocked them. For every filtered URL we determine the fraction of poisoned resolver blocking it. *Consistency* is the average of these fractions.

In MTNL, we found a total of 448 resolvers, out of which 383 were poisoned, *i.e.,* coverage was around 77%. Whereas, in BSNL we found only 17 poisoned resolvers out of a total of 182 (a smaller coverage of around 9.3%).

The consistency of each ISP can be inferred from Figure 4.2. Websites which are blocked in any of the two ISPs are represented on the X-axis. The percentage of resolvers blocking the website are represented on Y-axis. For the sake of preserving anonymity, we represent the sites

with unique numbers, rather than actual names. It can be clearly seen from the figure that in general a single website (*e.g.,* website ID 450) is blocked by more number of resolvers in MTNL (44%) than in BSNL (6.6%). The consistency metric in MTNL (42.4%) is also higher than that of BSNL (7.5%).

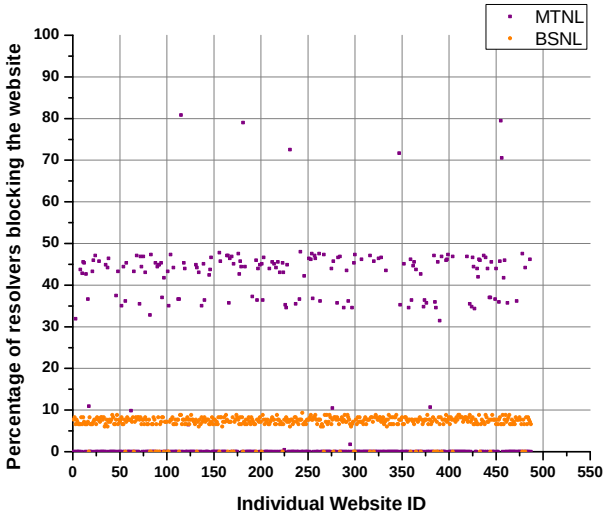

Figure 4.2: Consistency of DNS resolvers.

4.5.2 HTTP Filtering

We found HTTP filtering in four out of nine ISPs. As already discussed in Section 4.3.4, ISPs have deployed middleboxes which inspect the packets between the client and blocked websites with an intent to censor traffic.

We began by identifying all those websites, among the 1200 PBWs, that were censored by the ISP. For instance in Airtel, we observed a total of 234 websites to be censored. The corresponding number for the remaining three ISPs is presented in the last column of table 4.2. Using the approach described in Subsection 4.3.4 we determined that censorship was triggered *solely* due to request and not the response.

Finally, we attempted to find the actual network location of the middleboxes with a variant of INT, involving crafted HTTP `GET` requests. However, we were unable to pinpoint the exact IP address of the middleboxes in most of our measurements because of anonymization by the ISP. We discuss this in detail in Section 4.7.

We now present the behavior of the different types of middleboxes, we identified in the wild and describe their censorship mechanisms in detail.

Types of Middleboxes:
In our experiments, we identified two kinds of middleboxes — *viz., Interceptive Middlebox (IM)* and *Wiretapping Middlebox (WM)*. IMs are akin to transparent proxies which intercept connections between the client and server and establish a new connection to the server. In our studies, we found IMs that intercept client–PBW request and respond with censorship notification messages, while dropping the actual requests.

The other, *i.e.,* WMs, involve a host that is connected to an active network element via a wiretap. It receives a copy of all the packets exchanged and inspects for requests that need to be censored. Thereafter, it crafts responses and sends it back to the censor, with appropriate TCP header bits to terminate existing connections.

However, the WMs cannot outpace the client–PBW traffic flow, as they work with a copy of the packets. Thus, they are not as effective in filtering every single request with real-time efficiency, compared to IMs. For WMs, roughly in 3 out of 10 attempts to access blocked website, the middleboxes were ineffective in censoring the content. Whereas, for IMs, all such attempts were unsuccessful.

Interceptive Middleboxes We used our variant of INT (explained in Subsection 4.3.4) to first obtain the location of middlebox in the network path intervening the client and a filtered site.

Thereafter, we sent crafted HTTP `GET` requests, bearing the censored domain in the `Host` field, with TTL values large enough to get past the network hop, corresponding to the middlebox. Regardless of further increments to this IP TTL value, we never observed the expected ICMP TTL Expired responses, but rather received the censorship notification messages, *indicating that the middleboxes might be intercepting and dropping these requests*.

In order to verify that IMs are triggered only for the blocked domains, we sent a crafted GET request where `Host` field bore a non-censored domain, while also iteratively increasing IP TTL values. Interestingly, we always received ICMP TTL Expired messages, even when TTL was

large enough for the packets to transit the middlebox. This confirms that IMs only inspect `Host` field of HTTP `Get` request.

We went a step ahead and selected an array of hosts we controlled in different networks[3] outside Indian ISPs. On these machines, we hosted an ordinary webserver. From our client, hosted in the ISP under test, we created TCP connections to these remote machines. The remote host simultaneously monitored its own traffic. The client sent crafted `GET` requests with `Host` field requesting a censored domain. The destination IP address, however, was that of the remote host. Upon traversing a censorious middlebox positioned on the network path in-between, the client receives a censorship notification-cum-disconnection packet, with TCP FIN bit set. The subsequent 4-way disconnection always timed out (very likely dropped by the middlebox). Finally, the client attempted terminating the connection by sendingR a RST packet.

The remote host however receives *none* of the packets, other than the initial handshake messages and a RST packet. But, the TCP sequence number of this RST packet differed from the terminal RST packet sent by the client, thereby confirming that it was sent by a middlebox. This confirmed the presence of IMs.

We repeated the same exercise, by replacing the `Host` field with that of an uncensored domain. Interestingly enough, the request reaches the remote host unfiltered. The functioning of the IMs can be schematically shown in Figure 4.3.

Wiretapping Middleboxes Similar to IMs, we used our own variant of INT, to first obtain the location of the middlebox in the network path intervening the client and a filtered site.

Thereafter, we sent HTTP `GET` request to a blocked domain. We then inspected the network traffic (at the client) for the said message exchanges through `pcap`, and observed that the client receives the censorship notification-cum-disconnection packet, with the forged IP address (of the server) and TCP FIN+PSH bits enabled, which thereby enforces connection termination. Further even before the termination process resolves, the client receives a fresh TCP RST packet from the middlebox, bearing the forged IP address of the server that forces the client to terminate

[3]Planetlab, cloud services and hosts in different universities

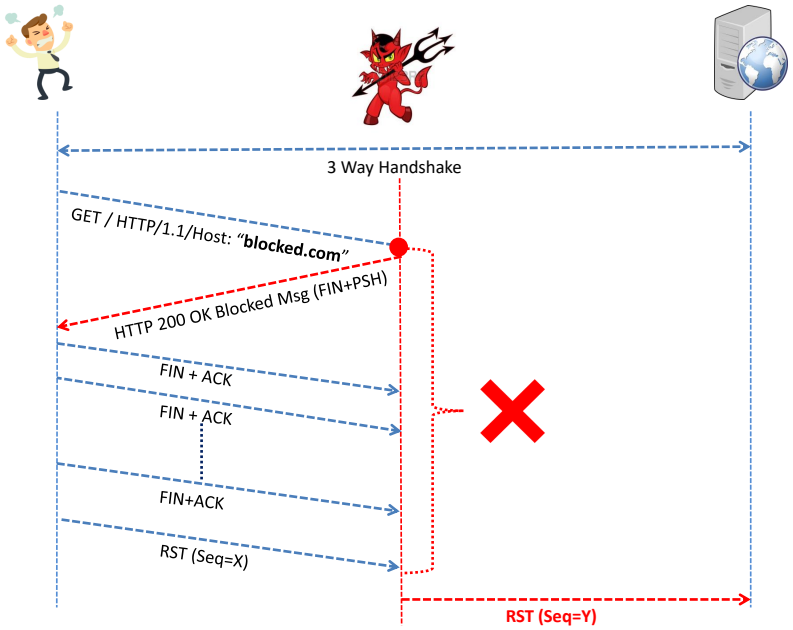

Figure 4.3: Censorship mechanism of a Interceptive Middlebox.

the connection immediately, regardless of whether the termination process, which is underway, completes or not.

Surprisingly, actual response from the filtered site eventually also arrived at the client, but the connection to the server was already terminated. The client responded with a TCP RST packet (as expected). *Such behavior indicate the presence of WMs.*

In order to confirm the censorship mechanism of WMs, we adopted an approach similar to the one described for IMs, involving remote servers under our control. We sent crafted HTTP `GET` requests bearing a filtered domain, to the remote servers under our control. These packets elicit the censorship notification-cum-disconnection messages, bearing the (forged) IP address of the remote hosts. The remote hosts, however, upon receiving the `GET` requests, ignore them as they do not host the requested domains. The behavior of WMs is shown in Figure 4.4.

Caveat: Are middleboxes stateful or do they inspect all packets? Our initial traffic inspections using `pcap` hint towards stateful middleboxes that commence traffic inspection only after complete TCP 3-way handshake is resolved.

To confirm our hunches we began with the client using `traceroute` command to record

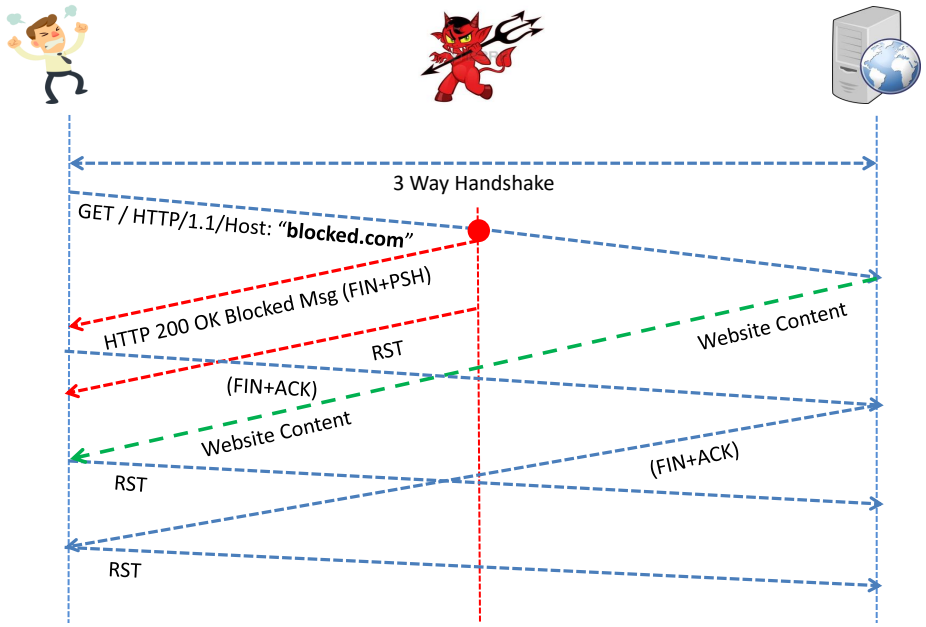

Figure 4.4: Censorship mechanism of a Wiretapping Middlebox.

the number of network hops between itself and the filtered site. Thereafter the client sends a TCP SYN packet with TTL just large enough to get the packet to the penultimate hop (and *not* the destination), thus avoiding a full-fledged TCP 3-way handshake.

Following this, the client sent a crafted `GET` request whose `Host` field pointed to a filtered domain bearing the TTL value, such that it expires upon reaching the penultimate hop. If the middleboxes commence traffic inspection upon observing every fresh TCP SYN packet, they must also then inspect the subsequent crafted `GET` request and respond back to the client with the censorship notification-cum-disconnection message. However, we never observed censorship in such cases. All other similar heuristics, such as starting by sending a SYN+ACK or not sending the final ACK of a regular 3-way handshake, but then sending the subsequent crafted `GET` request never elicited censorship messages.

Finally, a crafted HTTP `GET` request, bearing censorious domain requests in the `Host` field, but with no preceding TCP handshake, also does not seem to trigger censorship.

This confirms that the middleboxes are statefull and commence traffic inspection only when they observe a complete TCP handshake. These seem different from what were observed by Wang *et al.* [121] who looked into the architecture of Chinese censorship infrastructure.

Analyzing the Extent of HTTP Filtering

In order to analyze the extent of HTTP filtering in an ISP we proposed variants of the two previous metrics, *viz. coverage* and *consistency*.

1. ***Coverage***: A censorious ISP which is willing to use HTTP filtering must typically deploy middleboxes in a manner such that they intercept all the router-level paths inside an ISP. *Coverage* is the fraction of all such router-level paths that are intercepted by middleboxes (we call them as *poisoned paths*).

2. ***Consistency***: This metric attempts to answer the question — "how uniformly does an ISP block content" Ideally same number of websites must be blocked on all the poisoned paths of the ISP. In such a case we say the ISP is 100% consistent. For every filtered URL we determine the fraction of *poisoned* paths blocking it — (paths that block a particular website)/(paths that block any website). *Consistency* is the average of these fractions.

In order to find consistency and coverage we started our experiments with single vantage point (VP) in the ISPs. As already discussed earlier, HTTP censorship middleboxes are agnostic to the destination IP addresses of the HTTP `GET` requests (as long as they appear to be a part of an existing TCP connection). We harness this behavior of middlebox to find their coverage and placement statistics.

VP within ISPs: For each of the nine ISPs under considerations, we establish TCP connections with Alexa top 1000 websites from the client machine and sent `GET` requests with `Host` fields pointing to all 1200 PBWs. Even if for single `GET` request we observed censorship, we considered that path to be *poisoned* by the middlebox.
For Reliance Jio ISP, we only observed 64 out of 1000 paths to be tainted with middlebox. This gave us the hint that maybe middleboxes are not placed optimally to intercept a large fraction of ISP paths.

VPs outside the ISPs: To further test our observations with more VPs, we used various hosts outside India, but under our control (PlanetLab nodes, cloud infrastructure, and few other hosts in various universities). Our aim was to find the maximum number of middleboxes and the

fraction of paths they intercept, inside an ISP.

For doing so, we began by scanning all live IP prefixes[4] for a particular ISP, and searched for hosts with open TCP port 80. Then we randomly sample two such IPs per prefix. We recorded the router-level path leading and the number of hops to each of these prefixes, from each vantage point, using `traceroute`.

We tailored our INT, targeting traces to each of these IPs (for all ISPs), where for each targeted host, we send 1200 HTTP `GET` requests, corresponding to each of the PBWs. Upon obtaining the censorship notification-cum-disconnection response for even a single site, we considered the corresponding network path to be *poisoned*.

We summarize our results in table 4.2. Column two and three represents coverage for an ISP from a single VP within ISP and multiple VPs outside of the ISP. Column four describes which type of middlebox (interceptive or wiretap) is deployed in the ISP and last column describes the total number of websites blocked out of 1200 PBW. It can be observed that Idea has highest coverage (90%) whereas Vodafone has very low coverage value (2.5%).

For Reliance Jio, we observed a very different behavior. While we saw a relatively low coverage of about 6.4% when searching for middleboxes from a vantage point positioned inside the network, we found no middleboxes when probed from the remote VPs to IPs belonging to the ISP (with open TCP port 80). There are two possible explanations for this. Firstly, the middleboxes may be sub-optimally positioned and thus the requests from the remote VPs are not intercepted. Alternatively, the middleboxes maybe filtering request not only on domain names but also for source IPs belonging to Jio network itself. Since we were unable to pinpoint the IP addresses of the middleboxes, we lacked the necessary information to further quantify our findings.

After finding the coverage of different ISPs, we now present the results obtained from computing consistency for each of them. In Figure 4.5, X-axis represents websites which are blocked in any of the three ISPs (Vodafone, Airtel and Idea). The percentage of ISP paths that block a particular website are represented on Y-axis. It is evident from the figure, that

[4]Live IP prefixes were obtained from CIDR report [143].

ISP	Coverage (%) (VP: within ISP)	Coverage (%) (VPs: outside ISP)	Middle-Box Type	No. of websites blocked
Airtel	75.2	54.2	WM	234
Idea	92	90	IM	338
Vodafone	11	2.5	IM	483
Jio	6.4	0	WM	200

Table 4.2: HTTP filtering in different ISPs.

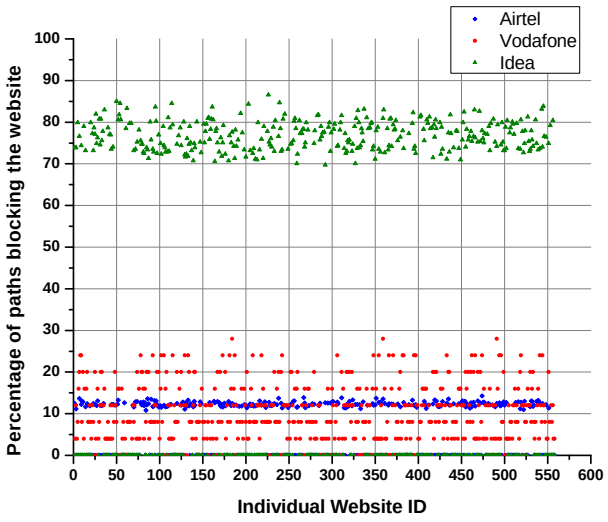

Figure 4.5: Consistency of middleboxes.

on an average Idea network has highest consistency (76.8%) followed by Airtel (12.3%) and then Vodafone (11.6%). One may conclude that in Idea network a single website is blocked on 76.8% of the poisoned paths, as opposed to Airtel and Vodafone in which it is blocked by only $\approx 11 - 12\%$ of the poisoned paths.

So far we discussed all details regarding HTTP filtering, but ignore HTTPS. We observed fewer than five instances of HTTPS filtering which were actually due to manipulated DNS responses by poisoned resolvers, and not because of SNI field in TLS client hello.

4.5.3 Filtering by Upstream Indian ISPs

An ISP shares contractual commercial agreements with its neighbors for routing Internet traffic among themselves [144]. In our study for ISPs like NKN, Sify[5] and Siti, we never observed any filtering caused by their own policies. Rather, all the censorship instances were solely due to the policies of its neighboring ISPs. Whereas, for MTNL and BSNL it is the cumulative effect of its own and neighbors' policies.

To precisely identify the locations of middleboxes, *i.e.,* which ISP they belong to, we used our tool INT. For cases, where we did not observe the IP address of the middleboxes, we used idiosyncrasies of different middleboxes to identify the ISPs they belongs to (as explained in Subsection 4.7.3).

ISPs where censorship was observed due to upstream providers, might indicate "collateral damage" (unintentional censorship) [13, 145] within the same country. This occurs when traffic of a non-censorious ISP is filtered due to a neighboring censorious ISP. In our studies we observed such unintentional censorship in several non-censorious ISPs. Table 4.3 summarizes our findings.

ISPs (cesnored)	**Neighboring ISPs** (causing censorship)
NKN	Vodafone (69), TATA (8)
Sify	TATA (142), Airtel (2)
Siti	Airtel (110)
MTNL	Airtel (25), TATA (134)
BSNL	Airtel (1), TATA (156)

Table 4.3: Filtering by upstream providers: Non-censorious ISPs observe censorship due to their censorious neighbors. *E.g.,* in NKN, we observed 69 websites were blocked by Vodafone and 8 were blocked by TATA communication.

[5]We could not independently study TATA communications because its customer business had closed during the course of study.

4.6 Anti-Censorship Approaches

We observed two types of censorship techniques in popular ISPs of India *viz.,* HTTP filtering and DNS poisoning. In order to bypass them, we opted for techniques relevant to the middleboxes involved. Our solutions are simple and extremely effective.

Evading *DNS Poisoning:* In order to circumvent poisoned DNS resolver, we tested using OpenDNS, Google's public DNS (8.8.8.8) and many other non-poisoned resolvers which belong to non-censorious countries like Ireland, Canada, and Sweden. With each of them, we were able to bypass the DNS based censorship.

Evading *HTTP Filtering:* As already explained in Section 4.3.4 middlebox gets triggered upon merely identifying a blocked domain in the `HOST` field of `GET` request. *Our goal is to craft a `GET` request, which is goes undetected by middlebox, but correctly interpreted by the actual website*. We tried various techniques involving string fudging [146], such as manipulating the `Host` field values, prepending `www` to the website name, changing cases of the keywords like `HTTP`, `GET` and `HOST`, adding spaces before and after the domain name *etc.*. Additionally we also tried approaches, like sending fragmented `GET` requests and using HTTP 2.0 as the underlying web protocol (instead of HTTP 1.1). Different approaches worked for subverting different middleboxes.

I. Wiretapping Middleboxes: There are two approaches with which we bypassed these middleboxes.

- Changing the case of `Host` keyword in the `GET` request: Most popular browsers, like Mozilla Firefox and Google Chrome, use the title case for the `Host` keyword. Merely changing the case (*e.g.* changing it to `HOst`, `HoST`, `HoSt` or `HOST`*etc.*) was sufficient for request to go undetected by the middleboxes (of Airtel and Jio), but resulting in response from the actual blocked webserver. This suggests that the webservers, corresponding to the PBWs, adhere to RFC 2616 [147] and accept the keyword `Host` agnostic of the case, while the middleboxes look for exact keyword matches.

- Dropping the packets with RST or FIN bit set: As mentioned earlier, the censorship

notification-cum-disconnection packet bears the TCP FIN bit set. Subsequently the middlebox also sends a TCP RST packet to enforce the client to disconnect.

Using `iptables` utility, all the packets (of blocked website's IP) which have FIN or RST bit set were dropped by the kernel. For Airtel, we observed that responses from middleboxes of Airtel always bear a fixed IP-Identifier value of 242. Thus, we added a general rule that FIN or RST packets with IP-Identifier field 242 must be dropped. This effectively filters the responses from the middleboxes.

Since the actual GET requests are not dropped by the middlebox, they reached the blocked website and elicit regular responses. These response containing the actual content of the website and are accepted by the client browser.

II. Interceptive Middleboxes: We further found two types of interceptive middleboxes *i.e.,* one which sends only censorship notification-cum-disconnection message to the client (*overt*) and other which sends only a RST packet to the client without any censorship notification (*covert*).

- *Overt Censorship*: To bypass such middleboxes which overtly censors the content, we fudged the `Host` field of the `GET` request. The standard domain request looks like "`Host`**:** blocked.com", *i.e.,* only one space between ':' and 'blocked.com'. But, instead, if we place additional spaces (or tab) in-between, *i.e.,* "`Host`**:** blocked.com", then the requests go undetected by the middleboxes, but servers interpret them correctly. Also, adding extra spaces (or tables) after the domain name works, *e.g.,* "`Host`**:** blocked.com ".
- *Covert Censorship*: For bypassing such middlebox we intentionally inserted multiple `Host` fields (with different domain names) embedded in the same `GET` request to check which of those is inspected by the middlebox. In all the cases, we observed that censorship is triggered upon inspecting *only* the last `Host` keyword. Thus, by appending an uncensored domain request to the array of such `Host` keywords, we were able to bypass the middleboxes, but the server also neglects it as the request is not a standard one. Thus we crafted an unusual `GET` request, which looked something like "GET / HTTP/1.1 `Host`**:**

blocked.com...\r\n\r\n `Host:` allowed.com". This request is neglected by the middlebox but on the other hand, accepted by the actual blocked website. Since middlebox is only looking at last `Host` keyword, it interprets that packet as non-suspicious and allows it to pass through. The server, on the other hand, treats the '\r\n\r\n' as the end of the GET request and the subsequent "`Host:` allowed.com" as a separate request. Thus, the client receives two responses from the website — the actual content due to the first `Host` field and the BAD REQUEST message for the subsequent one.

The Ant-Censorship tool — ESCAPER

We packaged all the aforementioned anti-censorship approaches into a tool called "ESCAPER". The tool runs a local HTTP proxy which needs to be configured in the browser. The client also needs to add a file mentioning the filtered websites it would like to access. The tool is written in Python and has been integrated with Mozilla Firefox browser. It runs on Windows and Linux platforms. Currently, it is maintained by us, and may be provided on-demand.

4.7 Discussion

4.7.1 Count of Middleboxes in the ISP

In the previous study on China [17] authors reported that they found 495 router interfaces that have filtering device attached to them. However, in India, we could not follow the same approach. Throughout our research, we used `traceroutes` and INT, with an intent of finding the location of censorship infrastructure. In all our tested ISPs, generally middlebox (or routers to which they are attached) show up as unresponsive routers (asterisked) when probed using `traceroute`. It is natural to ask if IP address of the middlebox is not known then can it be confirmed that the observed the censorship is because of the tested ISP or one of its upstream provider? We applied a few heuristics:

(1) On the paths where we observed IPs of middleboxes, first we confirmed that they belong to same tested ISP. Thereafter, we recorded the corresponding censorship notifications, and used

them to classify other anonymized middleboxes.

(2) In path segments where asterisked router appeared between visible ones, we checked if the latter belonged to the same ISP under test. If so then we assumed that anonymized IPs belong to the same ISP.

(3) The censorship notification messages have unique characteristics *e.g.,* in Airtel, the censorship notification packet has an embedded `iframe` which redirects to *"airtel.com/dot"* and in Reliance JIO censored response redirects to its own unique IP address. Using such unique characteristics we can easily identified the ISP of anonymized middlebox.

4.7.2 Issues with OONI

As already explained in Section 4.3.1, OONI performs two sets of experiments for a given list of PBWs (1) accessing sites from client machine and (2) and accessing the same from a control sever (of OONI). If discrepancies in IP address resolutions (DNS censorship) or retrieved site contents (HTTP censorship) are observed, OONI flags the PBW as censored.

However, we found that results of OONI were misleading. They suffer from both false positives and false negatives. We now outline few possible reasons for false positives (incorrect flagging of sites as being censored):

- An unavailable website, previously hosted on hosting services (like GoDaddy) if removed, may result in different HTTP responses when accessed from different locations — an artifact of distributed hosting. Though not a case of censorship, OONI flags them as filtered.
- Many websites have dynamic content such as live news feeds and advertisement embedded in the HTTP 200 OK messages that are often location dependent. These are also misclassified by OONI as being censored.

Also, OONI tool inspects differences in HTTP headers and body lengths of the response. If differences are greater than a threshold, it considers the site to be filtered. We observe

that for a website hosted on Content Distribution Network (CDN), the response at different geographic locations may come through different servers having obvious differences in the response metadata. In reality, such sites may not be blocked.

Thus when we created our scripts for detecting censorship, we only calculated the difference in the content of the response, and not the headers. If the difference is greater than the threshold, rather than directly reporting them as blocked, we manually verified them for blocking.

We now discuss why OONI often fails to detect a censored site (false negatives). In order to identify censorship, OONI calculates the differences between (1) lengths of HTTP responses (2) the HTTP header field names (3) the HTML title (obtained via the control server and directly through the ISP under test). A website is classified as censored, only when, the difference exists in all the three conditions.

Even if one of the aforementioned condition does not hold true [148, 140], OONI considers the website to be non-censored. The following are a few possible cases where OONI reports false negatives:

- We observed that for some websites, the response does not bear any content, rather a redirection link sent by the actual server. Similarly, in the censorship notification-cum-disconnection packets, there is an embedded `iframe` (which redirects to blocked page). For both the cases, the difference in the body length (of the responses) may be very less[6]. Thus, violating the first condition.
- OONI flags a website as non-censored if the header fields (and not their values) of both the HTTP response matches exactly[7]. In our measurements, we observed that most of the middleboxes use the same HTTP header as that of regular websites. Thus, the headers of censorship notification-cum-disconnection packets (generated by middleboxes) very often match the headers of the responses from the actual websites. So, OONI mistakenly classifies a censored website to be non-censored, thus violating the second condition.

Unlike the regular responses from actual websites, the censorship notification packets bore

[6] Other variants of such scenarios are also possible *e.g.,* a small sign up/login page upon accessing the website.
[7] As verified from source code.

no HTML tags. OONI compares the title tags *only if* atleast one word, in both the tags, is atleast five characters long. Thus, in the absence of the title tags, OONI ignores the inspection of the censorship notifications, thereby reporting incorrect results.

4.7.3 Idiosyncrasy of Middleboxes

- Idea middleboxes inspect traffic agnostic of their port number, while all the rest inspect only requests destined to TCP port 80.
- WM specific to Airtel have a unique characteristic — all packets generated from these middleboxes have a fixed IP-ID value (242) in the IP header; for all others', this is variable.
- Some otherwise unavailable websites[8] were still blocked by the ISPs (both through HTTP and DNS filtering). This implies that ISPs are not updating their blacklists.
- Middleboxes (IM and WM) maintain a state for all transiting TCP connections. They inspect all the connections for a duration of $2-3$ minutes, waiting for sensitive content to arrive. If they do not receive any packet in that duration, they time out and purge the corresponding TCP state data. However, if fresh packets (corresponding to individual flows, regardless of whether they carry `GET` requests or not) arrive with the $2-3$ minute window, the middleboxes reinstate the inspection timeout.

4.8 Concluding Remarks

In this work, we report a comprehensive analysis of censorship mechanism and infrastructure in nine popular ISPs of India. We commenced our research using popular censorship detection tool, OONI. However, since we observed high false positives and negatives, we discontinued using it. We developed our own automated approach (Interative Network Tracing), along with various heuristics, which we used to determine the type of censorship mechanism involved (and in some cases the approximate location of the censorship infrastructure as well). At every step

[8] Tested via Tor circuits ending in non-censorious country.

we confirm our findings against the ground truth, *an effort largely ignored by several others in the recent and distant past.*

We found DNS and HTTP filtering as the *only* techniques of censorship employed by these ISPs. Further, we evolved metrics, *viz. coverage* and *consistency* that respectively describe how well the censorship infrastructure covers the ISP and how consistent they are in censoring filtered domains. In passing, we also observed interesting cases of collateral damage within the ISPs of the same country. Finally, we developed novel anti-censorship techniques, involving local firewalling and manipulating the HTTP `GET` requests, through which we were able to bypass all forms of censorship without relying on conventional methods like proxies and VPNs.

Chapter 5

The Devil's in The Details: Placing Decoy Routers in the Internet

The government doesn't want any system of transmitting information to remain unbroken, unless it's under its own control.

Isaac Asimov

5.1 Introduction

Anti-censorship systems such as proxies or `Tor` [80] suffer from a double bind. To be useful, the entry point to the service must be discoverable to the user — typically, the citizen of a censorious country. On the other hand, as soon as the entry point becomes common knowledge, it also comes to the attention of the censoring government, who shuts it down [95]. *Decoy Routing*, a new anti-censorship paradigm [149, 150, 151, 152, 153, 10], attempts to disrupt this dynamic by using special routers as proxies, rather than end hosts. A decoy router (DR) lies on the path of traffic between the user inside a censorious country and an apparent ("overt") destination; when it senses secret handshake data embedded in the user's packets, it intercepts the packets and re-sends the message they carry to the real ("covert") destination. Note that the DR, being outside the censorious country, can freely communicate with the covert destination — and unlike

an end-host proxy, cannot *easily* be blacklisted.

However, "easy" is a relative term. In their paper on "Routing around Decoys" [154], Schuchard *et al.* propose that a sufficiently powerful adversary can simply route around ASes in the Internet where DRs are positioned. Houmansadr *et al.* [51] retort that such a move is extremely expensive, and in any case one could leave the adversary with no such option, *e.g.,* by placing DRs in enough ASes to completely encircle a censorious country. They then follow up with a model [155], where they frame the problem of placing DRs, versus the problem of bypassing them, as an adversarial game. *But the problem remains that the best known solutions still require the collaboration of several hundred ASes, in order to leave a* **single** *well-connected country*[1] *with no choice but to route through one of them. Further, such solutions require the computation of separate sets of ASes for each adversary nation [51, 155].*

Our first contribution in this research is a new approach to the question of placing DRs. In Decoy Routing, the router *intercepts* messages, from the user inside a censorious nation, en route to an overt destination. What if, instead of trying to intercept *all* the flows from a censorious country, we consider only the flows to the overt destinations? The overt destination is most likely a well-known site, often visited by citizens of the target country. [If not, it is very hard for users to discover; and when it is found, the sudden surge of traffic from users in China to some obscure website in *e.g.,* Turkmenistan will itself make the censor suspicious.]

As a first step, we started with the assumption that the overt destinations are popular sites (such as the Alexa top-10). We constructed a map of AS-level paths, connecting all ASes of the Internet to these, using the approach described by Gao *et al.* [21] (involving real BGP routing tables and inter-AS relationships [49]). We then identified the *"key" ASes* – those which appear most frequently on a large fraction of the paths. We find that ≈ 30 ASes appear in more than 90% of the paths to our target sites.

Our approach in this first step is not general; clearly, the adversary could block access to the entire Alexa top-10, to prevent users' traffic from reaching the DRs. So our second step is a study of how the "hardness" of the problem – finding the set of ASes that intercept several AS

[1] A "well-connected" country does not just refer to major powers like China; even *e.g.,* Venezuela is well-connected in this regard.

paths – varies as we change the set of possible overt destinations to the top-10, 20, 30, 50, 100, or 200 web sites.

Interestingly, we found the *same* set of 30 ASes intercept over 90% of paths in all cases – whether we consider paths leading to the top-10, 20 ... or 200. However, this result is easy to explain in hindsight. The Internet consists of ASes linked by peer-peer and provider-customer relationships; the "top of the hierarchy" or "core" consists of a few large multi-national ASes that peer with one another, and provide Internet access to most other ASes [14, 156, 157, 158]. Given such findings, and our experimental results with real paths, we come to a very powerful conclusion: only 30 ASes, all in non-censorious countries, are sufficient for a DR infrastructure that intercepts more than 90% of paths to important websites *in general*. In such a case, besides the reduction in the number of ASes compared to current solutions (about 30 times!) *this method needs to be run only once, rather than separately for each censorious country*. Our further experiments indicate that this is indeed the case – the power of these ASes is not limited to the top-100, they intercept over 90% of paths for other destinations as well. For example, with nine case studies of censorious countries, we found that these key ASes also intercept over 90% of the paths to 450 websites that are popular across these nations and also hosted outside their respective network boundaries.

On further analysis our AS-level results suggests that censorious countries in the Internet are less able to "route around decoys" than previously thought. About 30 ASes – 0.055% of the world ASes – intercept over 90% of paths to popular websites, and in particular, 99% of the paths originating from China. Our analysis also reveals that if censorious regimes choose to filter traffic along paths traversing the key ASes, they affect customers outside their network boundaries, and the extent of this "collateral damage" can be extremely high; *e.g.,* over 92% of all the network paths that traverse Chinese ASes originate beyond its network boundaries. We describe this in detail in Section 5.6.

For our second contribution in this research, we raise a new question. DR placement is not limited to AS selection! A large AS has thousands of routers; where exactly in the AS should DRs be placed? In this first study that uses intra-AS mapping (*viz., Rocketfuel* [79]) to answer the aforementioned question, we find that while the number of ASes required for a world-wide

DR framework is very small (30), we need to replace on average 400 routers per AS with DRs.

We conclude that, while a global DR system may involve only 30 ASes, a practical one would still require placing over 11,700 DRs in about 13 different countries. In fact, the problem remains challenging even if we provide Decoy Routing to citizens of a single country: against a very weak adversary, Syria (contained by only 3 ASes), a DR framework would involve 1,117 routers. *No existing DR architectures have been shown to process requests at line rates of network backbone routers*[2]*, nor has an implementation on existing high-speed routers been developed.* Unless we can deploy Decoy Routing on existing (or augmented) networking infrastructure, and can handle the high speeds, we will need to replace infrastructure at costs of over ten billion dollars (for example, for Level-3 Communications alone, *i.e.,* AS 3356 and AS 3549, the cost is 1.4 billion USD at Parulkar and McKeown's [160] estimate of 885,000 USD/ router), *plus* implementation, downtime, and debugging costs.

5.2 Background and Related Research

This section presents the relevant background for our work, and a brief discussion of how it fits into the existing literature.

5.2.1 Network Anti-censorship and Decoy Routing

The general area of our work is the use of proxy servers to circumvent censorship. Popular anti-censorship solutions, such as Tor [80][3], are no longer powerful enough when the adversary is a sophisticated nation-state: there exist techniques to detect TLS flows carrying Tor [161, 162]. More generally, traffic for most proxy based solutions can be detected and censored [163, 164], even if camouflaged [165].

Decoy Routing [149] takes a new direction where proxying is performed by special network routers called *Decoy Routers*. We sketch the basic mechanism in brief.

[2] Which is of the order of Tbps [159]

[3] Onion routing was originally designed to ensure anonymity over the Internet, but as it tunnels encrypted messages through a distributed network of proxies, it is also suitable for evading censorship.

- The user of Decoy Routing is hosted within a censorious ISP network, but wishes to communicate with network destinations censored by its ISP. To achieve this, it sends packets addressed to an innocuous-looking website, known as the *overt destination*. (The packets are encrypted using TLS, so the ISP cannot see the contents, and the header shows that they are meant for an unfiltered destination.)
- These innocuous-appearing packets, allowed out of the censoring ISP, carry a small, steganographic message, usually encoded in the protocol headers.
- On their way to the overt destination, if the packets pass through a DR, the steganographic message acts as a secret handshake. Instead of forwarding them, the DR decrypts their payload (the key, the TLS shared secret, is also sent as part of the secret message); establishes a new connection to the filtered site — the true, *covert destination*; and sends the payload to this covert destination.

Thus, a DR acts as a proxy, covertly communicating with a blocked site on behalf of the user. This procedure, *end-to-middle (E2M)* censorship circumvention, is shown in Figure 5.1. Actual implementations of Decoy Routing – Telex [151], Cirripede [150], TapDance [152], Rebound [153] and Slitheen [10] – have different features (message replay protection, tolerance of asymmetry in routing, inline blocking of traffic to/from overt destination, implementation of secret handshake, *etc.*), but share the basic design outlined above. This design decision stems from the realization that it is much harder for the censor to prevent the packets passing through a router, than it is to block an end host. But *how* hard it really is for the censor to circumvent DRs, and *where* the routers should be placed, is an active research question, as we discuss in the next Subsection.

5.2.2 On The Placement of DRs

Where should DRs be placed in the Internet? This question was first raised by the Cirripede project [150], where the authors claim that (against an adversary who is ignorant of Decoy Routing), placing DRs in just two tier-1 ASes is sufficient to serve all clients worldwide.

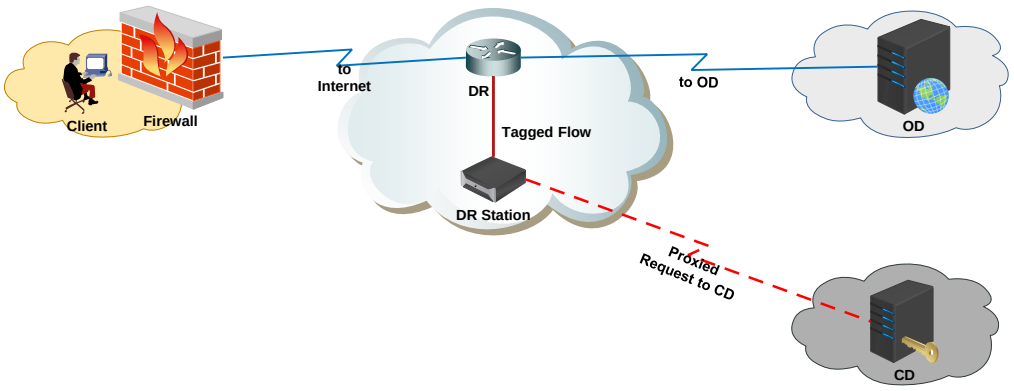

Figure 5.1: Decoy Routing in Action: Clients in a censorious ISP bypass the filter by sending packets apparently addressed to a non-filtered overt destination (OD). En route, the packets traverse a DR, which sees the secret message; identifies them for special handling; decrypts them; and sends their payload to the real, covert destination (CD).

The next major step was by Schuchard *et al.* [154], who argue that a powerful adversary *e.g.,* China will eventually figure out which ASes have DRs in them, and can then simply redirect its traffic to avoid these ASes — the Routing Around Decoys (RAD) attack. The authors map the Internet at AS level (ASes and their connections) to show that government-level adversaries (China, Iran, *etc.*) have connections to many ASes, and thus enough alternative paths to route around a particular AS. Specifically, avoiding the top 100 ASes (by degree in the CAIDA connectivity graph) disconnects China from only 2.3% of web destinations. Kim *et al.* [166] also suggest a graph theoretic approach to solve the problem, involving hypothetical network graphs. Thus, so far, the authors consider hypothetical Internet topologies, without considering how network routes are determined by *inter-AS relationships* [167].

One may argue that the *adversary may force user traffic to be routed around DRs hosted in such prominent ASes*. Houmansadr *et al.* [51], challenging Schuchard*et al.*'s effort, counter such arguments, by showing that the actual costs of RAD are too high to be borne. They also question Schuchard *et al.*'s assumption that DRs may be placed in randomly-chosen ASes. 86.2% of the ASes on the Internet are *origin* ASes (*i.e.,* they do not transport traffic of other ASes); a random placement mostly chooses such ASes, and it is possible to do much better if the ASes are chosen strategically. The authors propose two ways to do this:

1. *Sorted placement*. ASes that appear most frequently in the *adversary's* routing tables.

2. *Strategic random.* ASes chosen randomly, but only among those ASes that have a large enough *customer-cone*[4].

But while this approach is certainly better than random, it does require that ASes be computed afresh for every adversary, *e.g.,* it needs ≈ 900 ASes *for China alone*.

The first contribution of this research is a new approach for placing DRs: we select the ASes that appear most frequently in paths from all ASes to popular websites (potential Overt Destinations), as candidates for placing DRs.

We tried diverse sets of accessed websites (popular globally as well as those in several censorious nations) as potential Overt Destinations and discovered that set of key ASes does not vary significantly. We suggest that these are very likely "heavy-hitter" ASes of the Internet, and thus good candidates for installing DRs. Our results indicate that *30 ASes suffice to provide Decoy Routing worldwide*; this holds true even when we restrict ourselves from using ASes in censorious countries such as Russia or China.

Unlike Houmansadr*et al.*'s approach, our approach does **not** require computing the candidate ASes for individual countries, *a one time estimation suffices.* Moreover, we correct some incorrect assumptions made by the earlier authors. Houmansadr [51] use customer-cone size as a metric to choose ASes, assuming that it is a good predictor of how many flows they carry; we explain in Section 5.6 and in the appendix why it is not.

The dominance of the heavy hitter AS, while surprising, is supported by earlier work — particularly, by the classical paper of Rexford *et al.* [14], which suggests that the Internet has a hierarchical structure rooted at a few "core" ASes. Our work demonstrates, using real BGP routes, that *with only about* 30 *ASes, we can ensure that the client can get access to DR anywhere in the world (with* 99% *probability of success in two attempts).*

Our second, and more important, contribution is to demonstrate that *even though the number of ASes needed for a DR infrastructure is small, the actual number of routers that are to be replaced with DRs is large.* We map ASes at the router-level, using Rocketfuel [79], and identify

[4]Customer cone refers to customers, customers of customers, etc. In other words, a selected AS must be a significant provider to other ASes.

the specific network elements that potentially need to be replaced by DRs; on average, for each AS we need to deploy several hundred DRs. We suggest that the cost of such "major surgery" effectively removes the possibility that ASes would operate such a project *pro bono*, and raises the question of how such an infrastructure may be economically feasible [155].

5.2.3 Mapping the Internet

Our work depends on finding the paths to a particular destination taken by Internet traffic. In this subsection, we give a short introduction to Internet mapping, and explain our method of mapping.

The Internet consists of routers and hosts, organized into networks called Autonomous Systems (ASes). These networks operate independently, but collaborate to route traffic among themselves. ASes can be customers, peers, or providers to other ASes; besides a physical connection, there must be an acceptable business relationship between two ASes, before they route traffic through each other[5].

The AS-level path between two hosts on the Internet is said to follow a "valley free" path [167], as the path first rises - an AS, then its provider, then a provider of the provider, etc.; peaks - or plateaus, as it crosses through several peering links - and then descends, through provider-to-customer links, until it reaches the destination. There are no "valleys" in the path; no provider-to-customer links between two customer-to-provider links, or vice versa.

Mapping the Internet involves two tasks — Finding inter-AS connections (and relationships) and mapping routers and hosts (and their connections, inside ASes).

AS-level mapping.: Projects such as CAIDA Ark [42] and *iPlane* [43] map Internet routes with `traceroute`. Traceroute returns router-level paths from a source to a destination, hop-by-hop; the map is built by running traceroute from distributed volunteer nodes to various $/24$ prefixes. This data is consolidated into a graph where the nodes represent ASes, and edges represent links between them.

[5] A customer AS routes traffic through its providers; but providers do not route "through" traffic through their customers. The only traffic a provider sends a customer, is meant for that customer, or *its* customers, and so on.

Such approaches are generally limited by the network locations and availability of the volunteer nodes; they may not provide the AS-level path between any two randomly chosen ASes, and even where they do, they may be inaccurate.

In our research we used the approach of Gao and Qiu [21], that uses RIBs collected from the Routeviews project [22] and "stitches together" known links, thus constructing paths to our target sites from every AS in the Internet. This approach has been used in the past by others [168, 169]. (The algorithm has been already explained in Chapter 2 Subsection 2.2.1).

Router-level mapping: A large AS, such as an ISP, generally has several thousand routers. In theory, it is possible to repeat our approach for inter-AS mapping (where we use BGP information), and map the internal structure of ASes using their SNMP Management Information Bases (MIBs) [170]. However, we have no access to this data. Instead, we mapped the routers in ASes of interest using the Rocketfuel approach [79] (this involves running `traceroute` probes from looking glass servers [171] to prefixes inside a chosen AS). Thereafter, *IP aliases*[6] are resolved using `Midar` [102].

5.3 Motivation

The problem in this research is to determine where in the Internet we should place DRs, in order to intercept large fraction of network paths. The current state of the art [51] chooses ASes which are strongly linked with each target country (therefore intercepting much of their traffic), and whose customer cone size exceeds some threshold. However, this approach has the following limitations:

1. New ASes must be identified for each adversary nation.
2. This set of ASes is quite large. ≈ 900 ASes for China, 850 for Venezuela, *etc.*
3. Customer-cone size does not seem to be an effective metric for choosing ASes that appear frequently in real routes (candidates for DR placement)[7].

[6] Different interfaces of the same router, with different IP addresses, are called IP aliases
[7] We mention the reasons in Section 5.6. Details are provided in the Appendix.

4. A large AS has thousands of routers, spread across several countries. Current methods identify the ASes to place DRs in - but not *where* in the AS they should be placed.

In order to address these limitations, we construct a map of the Internet, and select the ASes that occur most frequently in our paths (estimated using real BGP routing tables), instead of any other metric. Next, we map these ASes to identify their key routers; this allows us to estimate the number of DRs we need to be able to intercept a large fraction of Internet traffic.

5.4 Methods: Data Collection and Algorithm

This section presents our algorithm for identifying key ASes in the Internet, and key routers in these ASes. Our focus in this section is on finding ASes and routers that intercept the paths from all ASes to important destinations (Alexa top-10, top-20 etc.) We also describe how we verify that our results are more general, *i.e.,* that our key ASes and routers also intercept paths to other destinations besides the top-n website. This is covered in more detail in Section 5.6. Our network mapping process consists of two phases[8].

- First, we build an AS-level Internet map, consisting of paths connecting popular websites and all the ASes of the Internet. We identify ASes that appear most frequently in those paths as key ASes (for hosting DRs).
- In the second phase, we estimate the router – level topology of key ASes to identify key routers – the actual routers inside the ASes that transport the majority of traffic.

5.4.1 Generating AS level maps

For the first phase of network mapping, we used the approach presented by Gao *et al.* [21]. AS paths are collected from BGP paths at Internet Exchange Points (IXes) [22]. These tables, however, do not contain paths originating at every AS; Gao *et al.*'s approach infers paths

[8]Our original plan was to map the entire Internet at the router level, and identify the key routers directly. Unfortunately, no existing techniques scale to mapping the Internet accurately at such fine granularity.

originating from every AS, using the existing BGP paths. ASes are appended to existing paths by selecting those that most frequently appear adjacent to ASes on the said BGP paths, without invalidating the path's *valley-free* property[9]. The aim is to build paths connecting every AS in the Internet to a given IP prefix. For our analysis, we used snapshots of BGP RIBs collected from 15 vantage points [22]. Our original approach involved choosing the top-10 most popular sites, finding the paths from all ASes to their corresponding prefixes, and identifying the *most frequently appearing ASes* on these paths.

As presented in Section 5.5, we found a small set of ASes that appear in more than 90% of the paths to these popular destinations from all ASes. We then increased the number of popular destinations – top-10, 30, 50, upto 100 – and estimated the paths to the corresponding prefixes from all ASes. At each step we identified the set of ASes which appear most frequently in the paths. As we show in Section 5.5, the rough set of ASes remained almost unchanged as we varied the number of destinations.

These results suggest that the ASes identified were likely "heavy-hitter" ASes of the Internet, potentially suitable candidates for DRs placement, as they may intercept large fraction of network paths, originating at ASes around the world. As a caveat, we know that the Internet has a hierarchical structure, rooted at a few "core ASes" which peer with one another and intercept large fraction of network routes [14]. But to test this claim, we had to answer two questions – (a) Was it necessary to select exactly, and all, these ASes for placing DRs? Some of these ASes were in censorious countries. (b) How could we validate that our observations, that a great majority of paths were intercepted by these ASes, are not limited to the target OD sites we studied?

In order to answer the first question, we investigated the impact of replacing our key ASes in Russia, China, *etc.*, with the next best choice: ASes ranked 31-50 by path frequency, but in non-censorious countries. We found the path coverage remained over 90%.

To answer the second question, we took the key ASes computed for the top-100 sites (say, Set-A). Next, taking sites ranked 101-200 on Alexa (Set-B), we computed the paths to these sites

[9]The AS-level path between two hosts on the Internet is said to follow a "valley free" path, as the path first rises - an AS, then its provider, then a provider of the provider, etc.; peaks - or plateaus, as it crosses through several peering links - and then descends, through provider-to-customer links, until it reaches the destination. There must be no no provider-to-customer links between two customer-to-provider links ("valleys")

for all ASes. We discovered that the key ASes, *computed using the paths for Set-A*, continue to intercept over 90% paths for Set-B.

Finally, we also computed the paths corresponding to the 50 most popular websites in each of nine different censorious nations (say, Set-C). The same key ASes also intercepted over 90% of the AS-level paths to destinations in Set-C.

Several months after our initial route collection, we repeated our experiments, and found the same set of key ASes intercepting over 90% of the paths to Set-A, Set-B, and Set-C.

Our approach differs from what was proposed previously [51, 154]. The authors either chose Tier–1 ASes, or those that had large customer-cone sizes. We show in Section 5.6, and in the Appendix, that customer-cone size is poorly correlated to the number of network paths that traverse an AS (path frequency)–*the latter being a better metric to select candidate ASes for DR placement.*

We note that Gao *et al.*'s algorithm generates the *request* paths (connecting all ASes *to* selected IP-prefixes), and not the *reply* paths (*from* the prefixes to the ASes). It is natural to ask whether asymmetry in routing might impact the strategy for placing DRs. However, the latest DR architectures, such as TapDance [152], are *agnostic* to path symmetricity[10]. This greatly simplifies the Decoy Router placement problem: we only need to place a DR on the path from the user to the OD, and not necessarily on the return path.

5.4.2 Creating Router Level Maps

After identifying key ASes in the Internet, as above, we were still left with the problem of *where* in the AS to put DRs. An AS involves a complex topology of routers and hosts; even the AS administrator, who knows the internal topology, may not know how frequently a router appears in actual network paths. When approaching AS admins to ask them to implement Decoy Routing, it is helpful to estimate how many (and which) routers they will need to replace. We therefore identified the actual routers that transport most of the ASes' traffic, using Rocketfuel [79] as

[10]Responses from the overt destination are suppressed by manipulating http protocol states, without requiring the DR's intervention.

follows:

- For each chosen AS, we identified the prefixes it advertises (from `cidr-report.org`). From 390 planetlab nodes, we targeted `Traceroute` probes to three representative IP addresses, corresponding to each prefix; we thus obtain router-level paths terminating at these prefixes. To capture paths transiting the said AS, we also ran `traceroute` probes targeted towards IP prefixes in its neighboring ASes.
- Using `Whois` [172], we inspected each `traceroute` paths to identify the first and last IP address belonging to the target AS. We denote these as the *edge routers* of an AS (as opposed to *core routers*, *i.e.,* the internal routers of the AS). We trim the traces down to the part between these edge routers, *i.e.,* inside the AS.
- The router IPs (belonging to the target AS), discovered through the above process, suffer from problems such as aliasing [90], so we resolved these aliases using the state-of-the-art alias resolution tool `Midar` [102].
- Finally, from the `traceroute` results we identified a minimum number of routers which cumulatively intercept over 90% of `traceroute` paths. When possible, we selected edge routers (as the edge routers cover 100% of paths through the AS). But in cases where some heavy-hitter edge and core routers intercept over 90% of paths, *and this set is smaller than the set of edge routers*, we selected those (the former) instead.

5.5 Data and Evaluations

5.5.1 Identification of Key ASes

As described in the previous section, we began by selecting a small set of globally popular websites (Alexa top-10), computed the AS-level paths to them, and identified ASes which appeared most frequently in these paths. We recomputed such paths by increasing the number of popular websites – top-30, 50, 70 and 100. As Figure 5.2 shows, the same number of roughly 30 ASes intercepted over 90% of the paths to these sites.

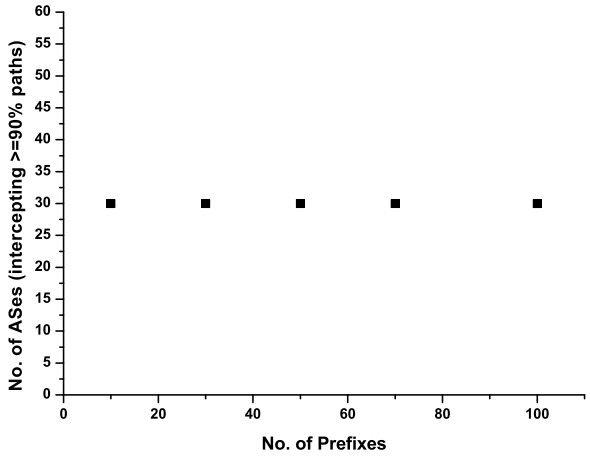

Figure 5.2: ASes needed to capture 90% of traffic paths to different sets of overt destinations (popular websites).

Figure 5.3 shows the CDF of the $4,497,547$ paths connecting the Alexa top-100 sites to all ASes, and their interception by our top-30 ASes. The X-axis represents the top-30 ASes ranked by their path frequencies; the Y-axis represents the actual fraction of paths. The highest-ranked AS, AS3356 (Level 3 Communications), intercepts $1,492,079$ paths (33.2% of all paths). The top 2 ASes, AS3356 and AS174 (Cogent Communications), intercept $2,028,831 (= 1,492,079 + 536,752)$ unique IP-prefix-to-AS paths, (45.1% of all paths). The top 30-ASes by path frequency together intercept 92.4% of all paths.

The details of these top-30 ASes are summarized in Table 1, which presents the corresponding ASes, their hosting country, and their ranks based on path frequency (P_{freq}) and customer-cone size (C_{size}). We highlight the ASes in countries known to censor Internet traffic, such as Russia and China [11].

Key ASes excluding censorious regimes

From Table 1, nine of our 30 key ASes lie in censorious countries and cannot host DRs. Finding alternatives is a legitimate concern. From the ASes ranked between $31-50$, we selected 9 new ASes headquartered in non-censorious regimes. This selection can be attributed to the fact that the ASes ranked $31-50$ are comparable to those ranked $11-30$ in terms of the paths

[11] As per censorship ratings by Freedom House Report [173] and the ONI [174].

ASN	Country	Rank (P_{freq})	Rank (C_{size})
3356	US	1	1
174	US	2	2
2914	US	3	5
1299	SE	4	4
3257	DE	5	3
6939	US	6	13
6461	US	7	8
6453	US	8	52
7018	US	9	17
10310	US	10	6
4134*	CN	11	10
3549	US	12	79
4837*	CN	13	85
209	US	14	19
9002	UA	15	97
6762*	IT	16	7
8359*	RU	17	22
2828	US	18	30
20485*	RU	19	21
16509	US	20	9
9498*	IN	21	18
4323	US	22	16
3216*	RU	23	99
2497	JP	24	15
701	US	25	12
12956	ES	26	65
37100	MU	27	23
4826*	AU	28	26
12389*	RU	29	67
1335	US	30	92

Table 5.1: Top 30 ASes that intercept more than 90% of paths. (ASes headquartered in censorious nations are highlighted.)

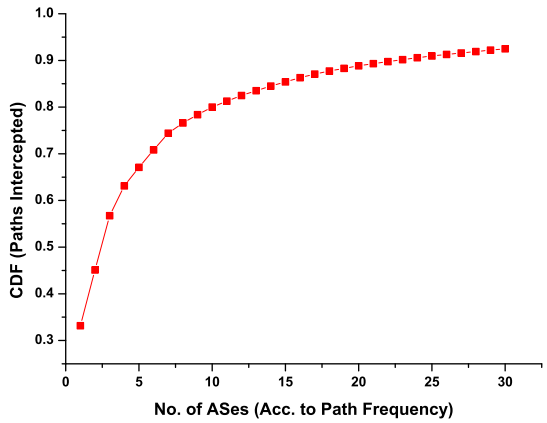

Figure 5.3: CDF: ASes and the fraction of paths they intercept. (CDFs are for paths to Alexa top-100 websites, unless otherwise stated).

they intersect. While the number of *unique* paths intercepted falls off rapidly, the *total* paths intercepted, *including overlaps*, does not (as seen in Figure 5.4).

We present our chosen replacement ASes in Table 5.2. Figure 5.5 presents the proportion of paths covered by key ASes in non-censorious nations alone (*i.e.,* redefining key AS to exclude ASes in censorious nations).

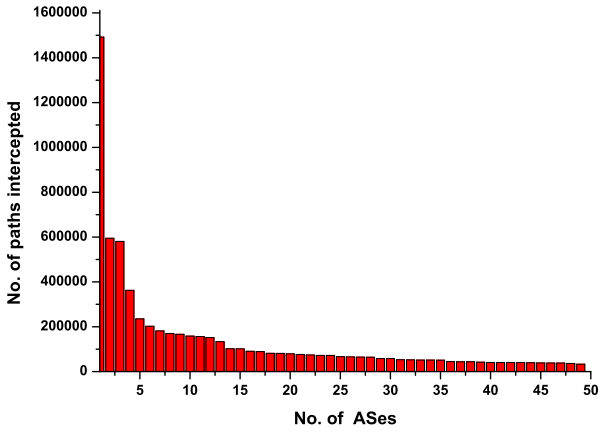

Figure 5.4: No. of paths intercepted by each of the top-50 ASes (sorted by path frequency).

We see that the cumulative path frequencies for the top-30 ASes hosted in non-censorious regimes (Figure 5.5) are similar to those seen for the top-30 ASes overall (Figure 5.3). In other words, the top-30 ASes *hosted in non-censorious regimes* are sufficient to intercept 90.23% of

ASN	Country	Rank (P_{freq})	Rank (C_{score})
13030	SW	31	84
1273	UK	32	83
16735	BZ	33	98
6830	EU	34	91
18881	BZ	35	95
3491	US	36	42
10026	HK	37	87
32787	US	39	93
1239	US	46	45

Table 5.2: ASes hosted in non-censorious nations ranked by path frequency (ranks >30 and <50)

the AS-to-prefix paths.

5.5.2 Identifying important routers inside key ASes

The second part of our research involves identifying the important routers inside key ASes. As described in Section 5.4.1, we used `Traceroute` to probe IPs in each prefix advertised by the key ASes. From these traces we determined the candidate routers that may be replaced with DRs.

We originally chose to näively replace edge routers with DRs, as these intercept all traffic entering and leaving an AS. However, we found that in many cases the total number of edge routers is significantly greater than the number of "heavy-hitter" routers – a set of edge and core routers that collectively appear in more than 90% of the `traceroute` paths for the AS.

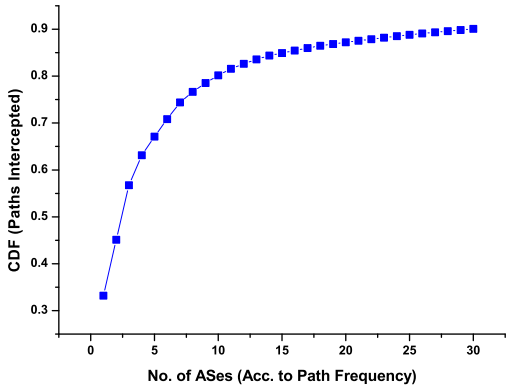

Figure 5.5: CDF of ASes (hosted in non-censorious ASes) according to fraction of paths that they intercept.

We therefore updated our approach. For each AS, we now find both sets (edge routers and heavy-hitter routers), and select the smaller set as the *key* routers, *i.e.,* the candidates for being replaced with DRs. For example, for AS 4134, we need only 179 heavy-hitter routers (including both edge and core routers) to capture over 90% of the paths, but 749 edge routers to intercept 100% paths, while for AS209 (Quest Communications), we choose the edge router set – about 1662 routers. We present our results in Table 5.3.

As mentioned previously, several of these ASes are hosted in censorious regimes, and so we identified the number of routers to be replaced with DRs in non-censorious countries. While the results presented in Table 5.3 represents the number of routers to be replaced for ASes presented in Table 5.1, Table 5.4 represent the number of routers for ASes in non-censorious nations presented in Table 5.2. The total number of routers that may be replaced across ASes in non-censorious ASes is $11,709$.

The $11,709$ key routers, across 30 key ASes that together intercept greater than 90% of network paths, together represent a formidable infrastructure, with equipment costs of 10.3 billion dollars. Converting these routers to DRs would involve massive system implementation, testing, deployment and related costs.

Implementation Details: Our AS-level map uses BGP Routing Information Base (RIB) data, which we obtain from 15 Internet Exchange points through routeviews [22], and AS relationship data from CAIDA [49]. The map was constructed using virtual machines with a total of 10 CPU cores (x64) and 24 GB RAM, running Ubuntu Linux (14.04LTS server). Our multi-threaded implementation of Gao's [21] algorithm took $\approx$ 3-4 hrs. to compute paths to 10 prefixes.

To identify key routers in an AS, we ran `traceroute` probes from 390 `planetlab` machines to three random IP's in each prefix advertised by the AS. Depending upon the number of prefixes advertised, and network latency, it took approximately $18 - 36$ hours to probe an AS, and $5 - 8$ hours for alias resolution.

5.6 Data Analysis and Discussion

Our method for placement of DRs, as presented in this research, has several major advantages:

1. The placement of DRs is *global*, and needs to be run *only once* to provide a small list of ASes that cover paths from all adversaries. (Existing approaches [51] require fresh

ASN	**# of Edge Routers (E)**	**# of Core Routers (C)**	**# of Heavy Hitter Routers (H)**	**# of DR's Required *min*(E, H)**
3356	707	303	576	576
174	165	1572	288	165
2914	134	2061	534	134
1299	493	1989	517	493
3257	762	2316	1483	762
6939	169	554	103	103
6461	105	850	45	45
6453	223	896	210	210
7018	359	6003	107	107
10310	161	156	106	106
4134	749	10078	177	177
3549	943	6227	5579	943
4837	1031	7350	2538	1031
209	1662	10842	8687	1662
9002	30	47	40	30
6762	154	333	238	154
8359	25	320	13	13
2828	116	1049	636	116
20485	506	206	193	193
16509	1244	5311	4644	1244
9498	320	199	269	269
4323	668	2548	2695	668
3216	305	1981	1769	305
2497	187	1078	133	133
701	1770	4417	2975	1770
12956	482	734	681	482
37100	14	72	59	14
4826	43	381	30	30
12389	322	2625	1898	322

Table 5.3: Edge routers, core routers, heavy-hitter routers and the routers required for replacement with DRs. Applying our router selection strategy, *e.g.,* for AS3356 – edge routers: 707 core routers: 303. Routers (both edge and core) covering 90% of the paths: 576. We thus select the latter. Total routers required for all the 30 ASes (headquartered in censorious and non-censorious nations) : $12,257$.

candidate ASes to place DRs for each adversary nation.)

2. The ASes selected are located far away from the adversary nations, and thus outside their geo-political and economic sphere of control. This makes it more difficult to bring pressure to bear on them.

3. The selected ASes lie on a very large fraction of paths. It is therefore hard for RAD adversaries [154] to bypass them without risking disconnection from all or most of the Internet.

One may ask why we only consider paths to the 100 most popular websites - what about paths to other IP prefixes? In this section, we explore such concerns, focusing on questions such as:

- Do our key ASes also intercept equally large fraction of paths to other unrelated sites (*e.g.*, less popular ones)?
- Particularly some users may consider completely different sites "popular" (*e.g.*, users in some countries may only be interested in sites available in their preferred language). Do our key ASes effectively cover paths to such sites?
- How important are the key ASes to actual censorious nations? If such nations chose to

ASN	**# of Edge Routers (E)**	**# of Core Routers (C)**	**# of Heavy Hitter Routers (H)**	**# of DR's Required *min*(E,H)**
13030	58	302	38	38
1273	156	1106	693	156
16735	12	43	37	12
6830	216	4048	1654	216
18881	338	3893	431	338
3491	698	1139	955	698
10026	170	765	346	170
32787	46	571	456	46
1239	242	1221	910	242

Table 5.4: Edge routers, core routers, heavy-hitters, and required DRs, for our "replacement" key ASes (from Table 5.2).

filter paths traversing these key ASes, how would it impact their downstream (foreign) customers?

Finally, we discuss the limitations of our method, and our plans for future work.

5.6.1 How general are our results?

Our data shows that a small fraction of ASes (≈ 30) cumulatively intercept over 90% of the total paths to popular web destinations (Alexa top 10, 20 ... 200)[12]. The question naturally follows whether these ASes are specific to the websites chosen for our study, or they intercept a similarly large fraction of all traffic on the Internet.

From Rexford's study of Internet structure [52], it is reasonable to deduce that there is indeed a small set of ASes - core ASes of the Internet - that cover a majority of routes of the Internet in general. So our task becomes, gathering evidence to show that our 30 key ASes intercept a large fraction of paths leading to various *other* destinations also.

To begin with, we estimated AS paths to Set B, the set of sites globally ranked 101–200 by Alexa. As Figure 5.6 shows, the 30 key ASes identified using paths to Set A, intercepted over 90% of the paths to Set B as well. The same pattern also holds for Set C - sites popular in censorious countries - discussed in the next subsection.

Finally, we repeated the entire experiment after a gap of four months. We again found the same 30 ASes intercepting over 90% of the paths (see Figure 4 in Appendix).

5.6.2 How Important are the Key ASes to Actual Adversarial Nations?

In order to answer this question, we began by measuring how well key ASes cover paths from individual adversary nations to globally important destinations. Our results, showing the fraction

[12]There is a third small caveat: Key ASes cover only more than 90%, and not 100% of the paths. But we now know, from Houmansadr [51], that it is not feasible for a country to launch a RAD attack and avoid 90% of paths. The only practical significance is that a user may not get a DR on their first attempt; but if a user probes for a DR, with greater than 99.9% probability she will succeed in three attempts (compared to the 30 attempts needed for earlier designs [150]), so this is not a major concern.

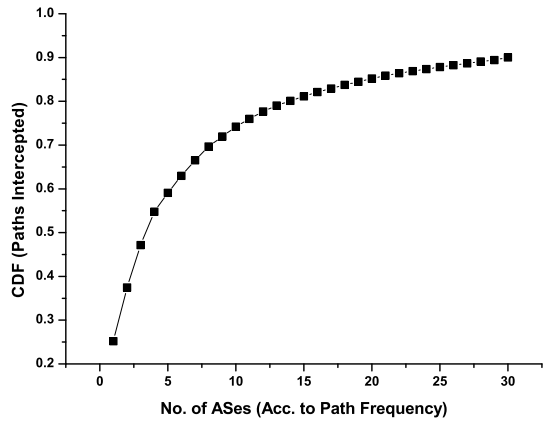

Figure 5.6: CDF of ASes according to fraction of paths they intercept (for Alexa top-101 to 200 websites).

of paths disconnected across 11 censorious nations, are presented in Figure 5.7. The horizontal axis has country names (as 2-letter initials); the vertical axis, the fraction of the paths covered by our key ASes. We see that, for example, our 30 key ASes cover 98.8% of paths from Chinese ASes to globally popular destinations, and at least 80% for nearly all adversary countries.

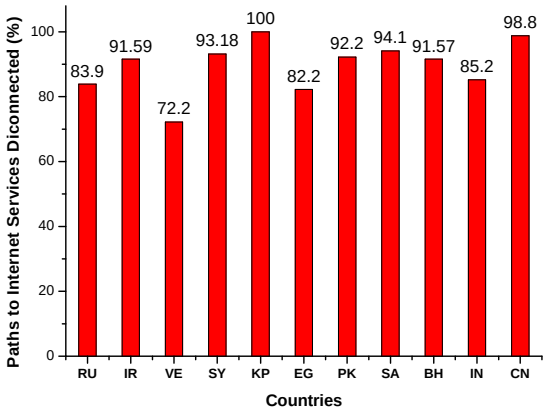

Figure 5.7: Eleven Censorious Nations: fractions of paths (to major websites) dependent on our 30 key ASes.

However, while these figures are encouraging, they are not enough. For some nations (*e.g.,* Iran or China), it might be argued that the loss of paths to globally important sites simply does not matter, as they have their own homegrown substitutes (*e.g.,* `facenama.com` and `renren.com` respectively for `facebook.com`).

In response to this concern, we investigated the popular web destinations in censorious

countries. As per Alexa [175], we find that these include not only local websites, *but also* several of the top-100 globally popular sites (search engines, social-media sites, cloud services, e-commerce sites *etc.*). In other words, while the choice of websites does vary across nations (*e.g.,* based on user's choice of language), web access is not as "insular" as one may fear.

For each of nine adversary countries studied by Verkamp [62] — China, Venezuela, Russia, Syria, Bahrain, Pakistan, Saudi Arabia, Egypt and Iran — we identified our Set C, consisting of the top 50 websites popular in each of these countries (and hosted outside their respective networks). Shown in Figure 5.8, our 30 key ASes intercept 93.3% of the paths originating or transiting these countries and leading to the sites in Set C.

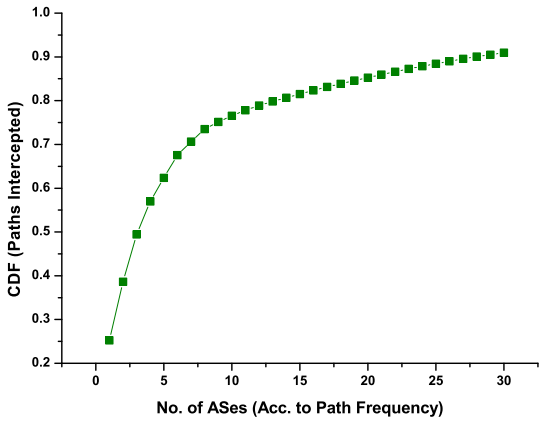

Figure 5.8: CDF of ASes according to fraction of paths intercepted (for websites popular in censorious nations).

In considering that different paths *originate in* and *transit through* a country, we further realized that avoiding key ASes might be expensive for a country in more ways than one – *collateral damage*.

Collateral damage: Collateral damage results when an AS filters sites, and also causes its customers to lose access [23]. If, for example, China was to boycott the paths routed through our chosen key ASes, Chinese people would lose access to much of the Internet (and certainly to most popular websites); but *so would customers of Chinese ASes*. It becomes a valid question to ask, how many customers are affected?

To answer this question, we inspected the paths through and from nine censorious countries.

Figure 5.9 shows the percentage of paths transiting censorious nations that originate at foreign ASes.

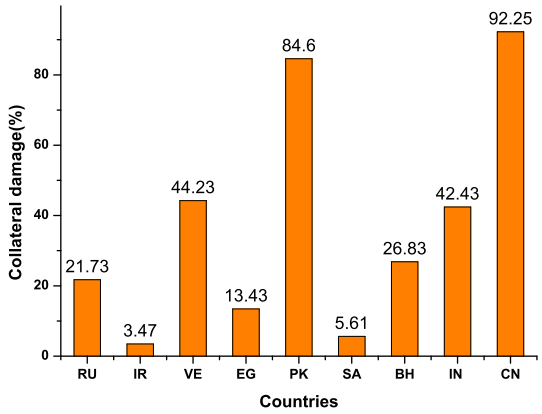

Figure 5.9: Collateral Damage: Percentage of paths transiting censorious nations that originate at foreign ASes.

We see that in the case of China, for example, filtering traffic through key ASes would affect a *very large* number of customers, over whom Chinese censorship policies have no control. $306,874$ AS paths, out of a total of $332,742$ paths involving Chinese ASes and leading to popular destinations — *i.e.,* 92.25% — *originate at ASes outside China*[13]. In fact, our data suggests that collateral damage to customers might be a way to put pressure on several censorious countries; we will explore this in future.

5.6.3 Might a Different Solution do Better?

The macroscopic analysis (of AS level topology) gives an impression that DR infrastructure is feasible, but the "devil in the details" is that the microscopic view (at router level) shows that we would need to convert thousands of routers into DRs. It is natural to ask whether this conclusion is just an artifact of our method, and whether an alternative approach might find a cheaper solution.

Our approach is not provably optimal. Indeed, we could get by with a smaller number of routers if we placed the DRs to intercept all traffic at a few, fixed overt destinations (*e.g.,*

[13] 362 particularly interesting paths originated at a Chinese AS, passed through non-Chinese ASes, then re-entered China and passed through Chinese ASes, before finally leaving for their destination.

Google). However, such a solution is fragile: the censor could simply filter traffic to these overt destinations. Our method of placing DRs uses far fewer ASes than any known comparable methods [51], and intercepts traffic to potential overt destinations (sites that are popular globally and also in censorious nations – for whom it matters most). Seeing how placing DRs in even our modest number of ASes is a major undertaking, we conclude that there is no "silver bullet" – *robust DR deployment is feasible, but implementing it is a serious challenge*.

5.6.4 Is it easier to cover single countries?

Our solution involves a single set of ASes that can serve as a DR framework for the overwhelming majority of traffic *globally*. We show that a global DR infrastructure is complex and likely expensive; but might it be feasible to target single censor countries?

We find that in case of major adversaries like China, the best solution is to use the same 30 ASes that we would use for a worldwide DR system. In case of some minor countries such as Syria (which has 2 ASes), Sorted Ring placement [51] does allow a simpler solution: we identified 3 ASes which intercept all Syrian AS level paths. But the router level maps of these ASes suggest that, even for Syria, we need $1,117$ DRs in 3 different ASes.

Our conclusion is that targeting a DR infrastructure to single countries is difficult even against relatively weak adversaries, and the best solution against strong adversaries (our solution in this research) is more expensive still.

5.6.5 How Economically Feasible is Decoy Routing?

Our results show that a comprehensive DR infrastructure would span about 30 ASes across ten countries, and require massive incentives. The question immediately arises whether existing business models for Decoy Routing [155], *i.e., central deployment* (where a single organization pays individual AS operators to deploy DRs) and *autonomous deployment* (where ASes individually deploy DRs and bill their users for usage), can reasonably provide such incentives.

In the case of central deployment, we note that unlike, for instance, Tor, this project will

depend on large-scale corporate participation. Tor is a globally distributed *volunteer* network, involving participants running the open-source Tor software on their (personal) end-hosts; the actual funding for the project only needs to support the developers, maintainers and some minimal infrastructure (*Directory Authority* servers, etc.) A worldwide DR framework needs to incentivize multiple multi-billion dollar companies to co-operate, and it is disturbingly likely that a single player who pays such incentives - whether a major company or a government - is motivated by its own agenda, rather than benevolence.

Autonomous deployment suffers from an even more serious issue. Decoy Routing obfuscates public knowledge of the deployment infrastructure (physical location and hosting network); *but such obfuscation also makes it difficult for users to target payment*. Any Internet based payment scheme would reveal the identities of DR hosting ASes to the clients, and in time, to their censorious ISPs. The adversary now simply blocks such Internet based payment transactions in order to prevent users from getting DR service; the whole "robust infrastructure that cannot be routed around" is rendered moot.

We therefore conclude that a practical DR infrastructure faces substantial challenges, and would likely only be possible with major support from one or more powerful nations.

5.6.6 Methods, Limitations, and Future Work.

This subsection is devoted to the choices we made, w.r.t. the design of our methods of network mapping. We explain our choices, their limitations, and how we propose to go forward in future.

Choice of AS: The first major question in our study, was how to choose key ASes. It may be argued that we could simply have chosen Tier-1 ASes – *i.e.,* ASes that have no provider – or ASes with the largest customer cone size (The customers, customers of customers, *etc.* of an AS are said to form its "customer cone".), based on publicly available data [49], as proposed earlier [51]. Why pick key ASes by path frequency? It turns out that there is indeed a good reason for directly choosing ASes by path frequency, *i.e.,* by how many of the paths they intercept.

A substantial fraction of AS paths traverse the customers of *"root"* ASes (*i.e.,* those with

very large customer cones) without traversing the root ASes themselves. *E.g.,* the traffic through AS9002 to AS2818 (`www.bbc.co.uk`) does not pass through AS3356, though it is the provider to both these ASes (see Figure 5.10).

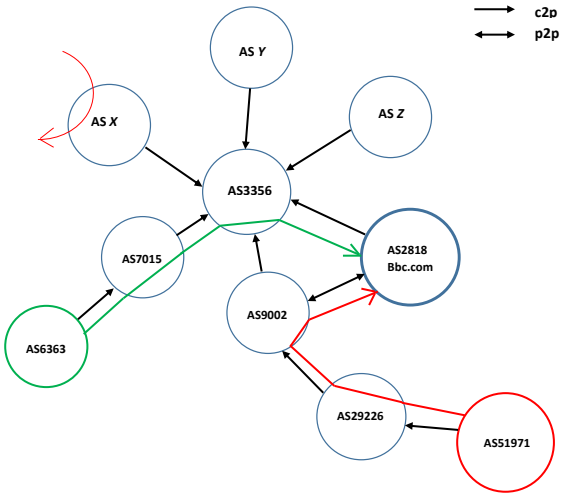

Figure 5.10: Valley free paths in the cone of AS3356. Green line: network path traversing AS3356 to reach AS2818 directly. Red lines: network path through one-hop customers of AS3356, but not AS3356 itself.

Unsurprisingly, 34.16% of the paths to top-100 IP-prefixes traverse the AS with the largest customer cone, AS3356 (cone size = $24,553$). But nearly as many paths, 33.17%, pass through its 1-hop (immediate) customers, and not AS3356. In the Appendix 7, we present more such figures (and Table 1), and show that Spearman's Rank Correlation coefficient [176] between AS ranks by path frequency and by cone size is only 0.2. Given the considerable fraction of paths which do not transit root ASes with large cone sizes (preferring to transit through their customer ASes instead), we conclude that *customer-cone size is not a good parameter to choose key ASes*.

AS path estimation: The two main methods of estimating an inter-AS topology are: (a) using `traceroute` traces (as in CAIDA Ark) (b) using BGP routing tables. Traceroute data, being constrained by the location of available probing nodes, is not sufficiently rich to estimate the actual path of traffic from every AS to a given prefix. Hence we choose the routing-table approach.

Previous efforts use simulated BGP paths [51], or paths derived from a Breadth-First traversal of inter-AS links [52]. We improve upon this by employing Gao's algorithm with *real* BGP tables (collected from various Internet Exchanges [22]), thereby estimating the *actual* paths from

a chosen IP prefix to all ASes (at a given point of time).

Of course, our map is still not perfect. As Gregori *et al.* [53, 54] show, publicly-available routing tables have biases, errors and bogus route advertisements. In order to address such issues, in future, we may cross-validate our map with different sources of data (and mapping algorithms).

Router level topology estimation: Router-level mapping of AS structure, uses `traceroute` probes from various `planetlab` hosts to IP addresses inside the ASes. We are limited by the fraction of routers discoverable through `traceroute` probes, and routers and middleboxes are sometimes set up to not respond to `ping` and `traceroute`. To limit this concern as far as possible, we use `Paris Traceroute`, with TCP probes. Secondly, we used `planetlab` nodes to launch `traceroute` probes, as the `looking-glass` servers (as used originally [79]) were unavailable at the time of our tests. There is always a chance some routes are simply not covered; increasing the number of probing hosts may improve our topology estimation, as new paths may be discovered by probing an IP address from different vantage points.

Further, in this work we identified border and internal routers (of an AS) using `traceroute` probes and WHOIS data (as already explained in Subsection 5.4.2). However, this approach is prone to errors. Thus, in future, we may resort to using advanced approaches like bdrmap [177], MAP-IT [178], and bdrmap-IT [179] (that has higher accuracy in identifying border routers of an AS).

5.7 Concluding Remarks

In this work, we have made two contributions towards answering the question of how to best place DRs in the Internet.

1. As our first contribution, we demonstrate that a small set of candidate ASes ($\approx$ 30) intercepts a very large fraction of paths (greater than 90%) to sites of interest, *i.e.,* potential overt destinations, *irrespective of the adversary country*. In other words, placing DRs in our ASes is sufficient to build a global DR framework. (As opposed to current approaches

[51, 155], which need the collaboration of over 800 ASes for a single adversary such as Venezuela or China.). We also observe that, if censorious regimes (like China) attempt to filter traffic along the paths transiting our 30 ASes, they will not only censor their own citizens, but many other residing outside their network boundaries (collateral damage).

2. Our second contribution is to explore the question of DR placement, not only at the AS-level, but at the router level. In practice, an AS is not a simple entity; it may have thousands of routers, and it is not obvious which of these should be replaced with DR. We find that, to intercept a large fraction of paths through an AS, we need a large number of both edge and core routers - typically several hundred (and in cases such as Quest Communications and Verizon, well over 1500 routers).

Thus, setting up a worldwide DR framework may require the collaboration of a small set of ASes (≈ 30). But even a single key AS, on an average, will need several hundred routers to intercept all the paths. We conclude that building a worldwide DR infrastructure is practically feasible, *but* ASes need sufficiently strong incentives to deploy a total of over $11,700$ DRs. We will explore such issues in our future work.

Chapter 6

Too Close for Comfort: Morasses of (Anti-) Censorship in the Era of CDNs

What we know is everything, it is our limit, of what we can be.

Julian Assange

6.1 Introduction

The Internet consists of more than 50,000 ASes that interact with one another through commercial business contracts (as *customers*, *peers*, or *providers* [167, 180]). In such arrangements, customer AS pays to a provider AS for Internet connectivity. Therefore, network researchers view the Internet as having a hierarchy — *i.e.,* a few ASes appear in a very large fraction of network paths [181, 182].

These few ASes are headquartered in only a handful of "powerful" countries (*e.g.,* the US) [145]. It is believed that since these nations intercept large fraction of network paths, they may be hegemonic over traffic originating from "underserved" nations. They likely surveil [9] (or censor [13]) transitory (or terminal) flows originating from these "underserved" nations. For instance, previous researchers [9] claimed that 95% of the Internet paths to Alexa top-1k

websites, for "undeserved" nations like India, either *transit or end* in "powerful" countries like US. Such observations hold good only *if* the client's Internet traffic crosses its national border.

However, popular web services rely on CDN infrastructure that ensures necessary redundancy (*e.g.,* Google Global Cache [183]) for providing high availability and performance [184]. Due to proliferation of CDNs, a country's traffic to popular destinations might not exit the national boundary. This may challenge the claim that powerful nation states have a hegemony over the Internet traffic. But, this may also inadvertently strengthen the ability of censorious regimes to coerce content providers for regulating access originating within their national boundaries [185]. Hence, in this research, we attempt to address some important concerns — *Are popular websites hosted within the same nation as that of the client? What impact does this have on the notion of nation state hegemony in the Internet?*

Unhindered web access using anti-censorship systems requires even more attention than earlier. Traditionally, such systems have been relying on publicly accessible proxies. But eventually adversaries discover such proxies and add them to their access blacklist. Recent attempts to disrupt this dynamic include Decoy Routing [149], Cache Browser [186], Meek [187], CovertCast [188] *etc.* Unlike regular proxies, these approaches rely on web content hosted outside the censor's boundary. Thus, we also answer questions like — *are these newer anti-censorship systems impacted by CDNs?*

Addressing such concerns requires identifying if popular websites (*e.g.* Alexa top-1k) are located within a nation's geographic boundary. Thus, we began by employing *Constraint Based Geolocation* (CBG) [189], a multilateration technique to geolocate Internet hosts. The process involves estimating the position of a target host using distance measurements from sufficient number of fixed reference points. These distances are estimated from RTTs between reference nodes and target hosts.

However, similar to others [190], we noted gross inaccuracies in CBG's geolocation — about 4000-5000 kms. Such errors often arise when the reference nodes and target hosts are far apart (at times even in different countries) [190]. We thus augment CBG, with novel heuristics, assuming the target host and reference nodes lie within same country under consideration. We call this

Region specific CBG (R-CBG). It decides whether a host is located in the considered country or not. R-CBG achieved $\approx 91\%$ accuracy when tested against the ground truth (domiciles of RIPE nodes obtained from RIPE Atlas project). Thereafter, we used R-CBG to geolocate Alexa top-1k websites. On an average, R-CBG reports about 80% of Alexa top-1k websites to be located in the same country as of the client.

This confirms that often web content is served to a client from caches located within its own country. Contrary to claims in [9], *traffic often terminates within the nations' boundaries*. Therefore, this makes ordinary Internet users more vulnerable to censorship by their respective countries, than to external surveillance.

Unfortunately, anti-censorship solutions do not alleviate this conundrum. They present a new double bind. Popular solutions, that rely on proxies are easily discovered and blocked by the adversary [191]. Futuristic solutions (*e.g.* Decoy Routing, CacheBrowser *etc.*) that do not bear trivially identifiable header signatures, rely on web requests to popular (unfiltered) sites hosted outside the censors' national boundary. Our measurements indicate that these solutions may also be severely impacted, as $\approx 80\%$ of Alexa top-1k sites are hosted within the client's nation. These sites either use anycast IPs or are hosted on non-CDN infrastructure.

Interestingly, we find that RTT by itself can be used to determine if a host is located inside a country or not. For targets outside a country, RTT is often greater than a threshold (and vice versa). We use this threshold to determine the relative position (w.r.t a country) of less popular sites (ranked above Alexa top-1k).

The following is the summary for our research efforts and findings:

- We quantify what fraction of Alexa top-1k websites are located within the same nation as that of client, for Saudi Arabia (SA), India (IN), Iran (IR), Brazil (BR) and United States (US). This involved using R-CBG, a heuristic driven multilateration technique, that compensates for location estimation errors. In $\approx 91\%$ of the cases R-CBG accurately judges if a target is hosted in a particular nation or not.

- R-CBG reports about 80% of Alexa top-1k websites to be located in the same country as

of the client. This has two major implications:

  - Contradiction of the earlier claims [9], that a few "powerful" nations may be hegemonic over majority of Internet traffic, particularly those that originate from "underserved" nations.

    Our findings challenge the the notion of nation state hegemony based on *transitory* Internet traffic only. Due to the presence of CDN front-ends, traffic originating from countries under test, does not exit the respective national boundaries. However, we acknowledge that powerful nations (where the CDNs are headquartered) might coerce the CDN providers to fetch data of clients (residing in other nations), from globally distributed CDN front-ends. We consider this beyond the scope of this work.

  - Hindrance towards adoptions of anti-censorship solutions (*e.g.* Decoy Routing, CacheBrowser *etc.*) which rely on popular web content, and require the traffic to leave the censor's boundary.

- For all five countries under consideration, we identify the type of CDN a website is using (Anycast [192] or DNS based [193]). We find that a large fraction of Alexa top-1k websites use Anycast CDNs (rather than DNS based). *E.g.*, in US, 59% of the Alexa websites use Anycast CDNs, 19% use DNS CDN and remaining 22% were hosted on non-CDN infrastructure.

- We observe that, by and large, RTT may itself be sufficient for identifying whether a website is located inside the nation or not. Our results reveal a clear distinction in RTT for websites that are hosted inside, versus those that are not. *E.g.*, in Iran for $> 99\%$ of the websites which were located inside, we observe RTT < 30 ms. Whereas, majority of those which were located outside had RTT > 100 ms. We further use RTT to classify 25,000 websites (Alexa top-5k for each of the five nations) as inside or outside.

6.2 Relevant Research

6.2.1 Proliferation of CDNs

Content Distribution Network (CDN) is a distributed architecture, that relies on *replica* servers to minimize end-users' access latency. It acts as an intermediary between the content publishers (site owners) and the end-users. Thus, it caches content at its edge servers (often called *front-ends* [194]).

In general, there are two types of CDNs — *Anycast* and *DNS* based CDNs [195, 196, 192]. In anycast based CDNs [192] (*e.g.,* Cloudflare), a single IP address is announced through multiple BGP advertisements, often from different geographic locations. A client's web request is directed to the closest possible front-end, based on BGP policies of the client's ISP.

However, in DNS based CDNs [193] (like Akamai), same website is resolved to different IP addresses, depending on the client's location. Whenever a client initiates a DNS query for a website, the resolver responds with an IP address which is (likely) closest to it. In general, DNS based CDNs maintain a separate mapping system to direct clients to their nearest front-ends.

In 2011 Ager *et al.* [197] conducted a measurement study to identify the extent of web content replication across different parts of the globe. They resolved Alexa popular websites, but (i) considered only DNS based CDNs in their study and (ii) relied on Maxmind database for geolocating the IP addresses which are known to be erroneous [198]. They reported that, for their tested domains, at least 46% of the popular domains were served from North America, 20% from Europe and 18% from Asia; the other three continents *viz.,* Africa, Oceania and South America did not serve much of the content. Additionally, when content was requested from North America 58.2% of the content was served from the same continent, while this number was only 26% for Asia.

To further study the proliferation of CDNs, many researchers focused on mapping the complex ecosystem of individual CDNs. For example, in 2013, Calder *et al.* [199] reported that Google had front-ends in over 100 countries and 768 ASes. Böttger *et al.* [200] studied Open Connect, the CDN owned by Netflix. Authors reported that IXPs play a vital role in large-scale

content delivery for Open Connects' world-wide customer base. Further, global CDNs have partnered with local CDNs of China to cater to their growing user base [201].

In 2018, Yeganeh *et al.* [184] studied the NetFlow data of a stub AS to understand the "locality of Internet traffic". They assumed RTT as distance metric for locality. Websites with low RTT were considered closer than those with higher. They reported that 90% of the traffic for the top 13 content providers was delivered within a 60ms RTT. Their results indicate that attempts made by different CDNs to bring content closer to the edge of the network are probably successful.

However, Scott *et al.* [202] took a different direction — rather than mapping a specific CDN, authors present the joint Analysis of CDNs and Internet censorship. They concluded that 20% of the Alexa top-10K websites were using CDNs. For the same set of websites, they found 4,819 instances of ISP level DNS hijacking in 117 countries.

6.2.2 Anti-Censorship Approaches

Free and open communication over the Internet, and its censorship, is a widely debated topic. There are numerous evidences of large scale Internet censorship [203, 204, 122, 205, 206] by various regimes. Thus, censorship circumvention systems have been devised [191, 207, 208]. Traditional systems rely on publicly accessible proxies. But eventually their IP addresses are discovered and blacklisted.

Systems like Decoy Routing, Cache Browser, Meek and CovertCast *etc.* are designed to avoid being discovered. All these approaches do not have any trivially identifiable protocol signatures (*e.g.* already known IP addresses). Rather, they primarily rely on web content hosted outside the adversaries' control. We now briefly explain them.

- ***Decoy Routing*** [149] employs routers (rather than end hosts) as proxies. Web requests carrying cryptographic signatures, sent to an apparent "overt" destination, are en route intercepted by the Decoy Router (DR) hosted beyond the censor's control. Based on the signatures, the DR identifies the packets and diverts them towards the intended "covert"

destination. DR requires the "overt" destination to be an unfiltered site, positioned outside the censor's control. This assumption may not always hold true if such sites are hosted on CDNs, located inside censorious countries. The Decoy Routed packets may never reach the DR.

- ***CacheBrowser*** [186] (and its successor CDNReaper [209]) leverages the fact, that it is extremely hard for an adversary to block all possible IP addresses associated with CDN hosted website. If client somehow learns an unblocked IP address (located outside the adversaries' control), it can access the website.
  However, if a website is anycasted, all its front-ends would use the same IP address. In the absence of numerous IP addresses (like in DNS based CDNs), it becomes trivial for the censor to blacklist the address. Moreover, anycasting often directs web requests to a front-end, likely in the same country (as the client).

- ***Meek*** [187] employs 'domain fronting' to by-pass the censorship. It assumes that censored and non-censored sites are both hosted on the same CDN. The apparent HTTPS communication of the client with a non-censored site (the front-end), bears a request to the blocked domain in the HTTP Host field. The front-end thus communicates with the blocked domain and responds with the filtered content. An eavesdropping censor sees nothing suspicious. However, when the front-ends are themselves located within the censor's control, their communication with blocked domains may be intercepted by the adversary. It is already known that significant fraction of web traffic exchanged between front-end and back-end (of the CDN) travels over public Internet [210].

- ***CovertCast*** [188] relies on sending the content of blocked websites via popular real-time encrypted video streaming services (*e.g.* YouTube). It is believed that censors are generally unwilling to block such services en masse. However, the authors acknowledge that the censors may coerce the service operators to block specific accounts associated with CovertCast. Interestingly, the presence of CDN hosted streaming sites located inside censorious nations, may facilitate such coercion.

6.2.3 Geolocation Techniques

Existing methods for geolocating Internet hosts are broadly classified as *active* and *passive*.

Passive geolocation methods mostly involve database lookups. These databases (*e.g.,* Maxmind [211]) are populated from various sources like reverse DNS lookups, RIRs, and ISPs [212]. Though widely used, these databases are notoriously erroneous [198, 213].

Active methods involve estimating geolocations from RTT measurements. The initial efforts [214] involved mapping nodes with unknown locations (also called *targets*), proximus to known locations. Later approaches mostly relied on Internet *multilateration* [189, 215, 216, 217, 218]. They involve estimating the position of a target host, using distance measurements from sufficient number of fixed reference points. The process involves plotting circles on the world-map with centers as the reference nodes and the distances as the radii. We call the region enclosed by such a circle as the *Probe Coverage Region (PCR)*. In theory multiple PCRs should intersect at exactly the target's location. However, due to measurement errors, rather than intersecting at one point, they often form an intersection area (ref. Figure 6.2). The target possibly lies in this intersection region.

There are various ways for estimating the distance a packet may travel. An example is the *Speed of Light (SOL)* constraint. It is largely believed that packets on network fiber cannot travel faster than two-thirds the speed of light (c) [189]. This principle provides an upper bound ($(2/3)c * OWD$, where OWD is one-way delay) on the estimated distance the packets may travel. These distances may be used for multilaterating a target.

Constraint Based Geolocation (CBG):

For geolocating Internet hosts, direct application of SOL constraint is not recommended. RTT (or OWD) is directly proportional to congestion (queuing delay). Even when geodesic distance between two hosts is small, RTT between them can be large (due to congestion). However, SOL constraint being completely ignorant of such factors, over-estimates the distance. Thus, Gueye *et al.* [189] proposed Constraint Based Geolocation (CBG), to incorporate these factors. It aims to

compute distances between reference nodes (in known locations) and the target host (which is to be geolocated), using RTT measurements.

Assuming the SOL constraint, CBG generates a "baseline" (ref. Figure 6.1) depicting the linear relationship between delays and distances, ignoring factors like queuing delays *etc.* Thus, SOL almost always over-estimates the actual distance a packet travels. Hence CBG, introduces a *calibration phase* to compensate for queuing delays which is ignored by the SOL constraint. In this phase, all the reference nodes `ping` each other and generate a scatterplot between delay (RTT) and distance (distance to other reference hosts is already known). From this plot a "bestline" is computed in a manner that: (1) All the data points (RTTs) are above this line and (2) This line is closest to all the data points in the plot (ref. Figure 6.1). Distance is estimated from RTT using the "bestline", instead of the "baseline." This distance is always expected to be less that what is estimated using the SOL constraint.

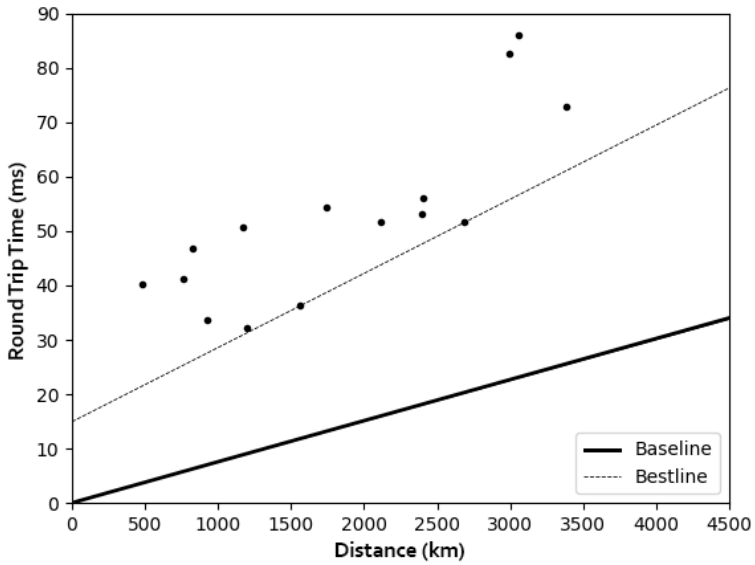

Figure 6.1: Baseline and Bestline for a RIPE probe.

With reference nodes as centres and these distances as radii, CBG creates PCRs on world map and computes the intersection region of PCRs. The target's location is assumed to be the centroid of this intersection region. Authors validated the locations predicted using CBG against the ground truth — nodes (with known locations) spread across US and Western Europe.

6.3 Methodology

Our research involved identifying what fraction of Alexa top-'n' websites reside within (or outside) the said nation using active geolocation techniques. Recently, Weinberg *et al.,* [190] empirically demonstrated that CBG outperforms other active methods like Octant [216] and Spotter [218]. Thus, we began by using CBG to geolocate Internet hosts. It is known that reference nodes far from the target do not contribute much to the geolocation process [215, 190]. Thus, using recommendations from previous efforts [190], we selected reference nodes in the same continent as the target. Our initial study for geolocating RIPE nodes as targets (with known locations) resulted in large errors — upto 4000 Kms. Similar errors (≈ 5000 Km) were also reported by Weinberg *et al.* [190]. To reduce errors in geolocation, we augment CBG by selecting reference nodes closer to the target (likely in the same country). We call this approach as *Region Specific CBG (R-CBG)* and used it to identify whether an IP address is positioned inside (or outside) the desired geographic area.

6.3.1 R-CBG: Improving Accuracy of CBG

We now explain how we improved the geolocation accuracy for CBG using R-CBG. Further, we also explain how it multilaterates anycasted IP addresses.

Our initial observation in geolocating RIPE probes with CBG, resulted in large errors, even when reference nodes were selected in the same continent as the target. Thus, we went a step ahead and individually selected the reference nodes, such that they were evenly distributed and *located either inside, or close to the geographic boundary*, of the nation under consideration. This vital step resulted in high accuracy of identifying whether the hosts, resides within the said nation or not.

As already mentioned in Subsection 6.2.3, to multilaterate an IP address, the reference nodes create probe coverage regions (PCRs) on the world map to produce an intersection region. The IP is expected to be located at the the centroid of this intersection of PCRs. In order to correctly identify, whether a node resides inside (or outside) the country, we present the following

four-point heuristic:

1. Intersection region of PCRs is completely within the boundary of the nation (ref. Figure 6.2).
2. Intersection region of PCRs cuts through the nation's boundary with centroid of region inside the nation.
3. Intersection region of PCRs cuts through the nation's boundary with centroid of the region outside the nation.
4. Intersection region of PCRs is very large and subsumes the entire nation's boundary (Fig. 6.3).

If condition (1) or (2) holds good, R-CBG predicts the target as an ***inside*** *node. Whereas if condition (3) or (4) holds good, R-CBG predicts the target as an* ***outside*** *node.*

For condition (1) and (2) to hold, the target should be located inside — either completely within or close to the boundary. However, in condition (3) although the target is close to the boundary, but is likely outside (in some neighboring country). Lastly, condition (4) represents the scenario where target is likely very far from the country. This is intuitive — assuming that reference nodes are located inside the country (*e.g.* India) and the target is located outside (*e.g.* SA), the radii computed will be large (and roughly same) for all the reference nodes, producing a large intersection area (shown in Figure 6.3).

To test our intuition, we use R-CBG for multilaterating the RIPE nodes. We tested it for five countries India (IN), Iran (IR), Saudi Arabia (SA), Brazil (BR) and United States (US). It resulted in high accuracy of about 91% in correctly disambiguating inside probes from the outside ones. We present the details in Subsection 6.4.1.

Multilaterating anycasted IP addresses (within a country): CDN providers like Cloudflare use *anycast* architecture. In anycast, a particular IP prefix is announced from multiple geographic

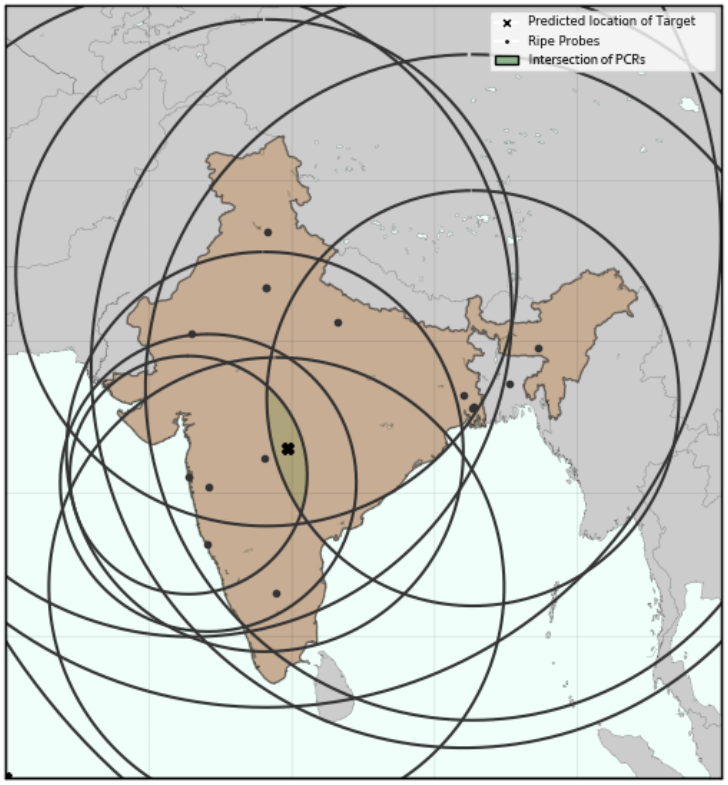

Figure 6.2: IP address located inside India. The intersection area is well contained within country.

locations (called as *anycasted sites*). This helps serve web content through redundant caches at various anycasted sites. Based on the BGP policies of the client's ISP, the web request is redirected to any one of the sites (likely the closest).

It is non-trivial to use native CBG, with globally distributed reference nodes, for identifying if an anycasted website is located *within a nation* or not. This is because each reference node ends up probing the same anycasted IP address, but at different sites. This results in multiple non-intersecting PCRs on the world map (ref. Figure 6.4). Since, CBG predicts the location of target as the centroid of the intersection region of PCRs, in this case it fails due to the lack of such a region.

However, for R-CBG, we select the reference nodes inside the nation under consideration. Thus, applying R-CBG to multilaterate an anycasted site might lead to majority of reference nodes pinging the same host. This would result in multiple PCRs forming an intersection region. Thereafter the four-point heuristic is applied to check if the target resides inside the nation or not.

In cases where a single IP address is anycasted at multiple sites *within the same country* (*e.g.*

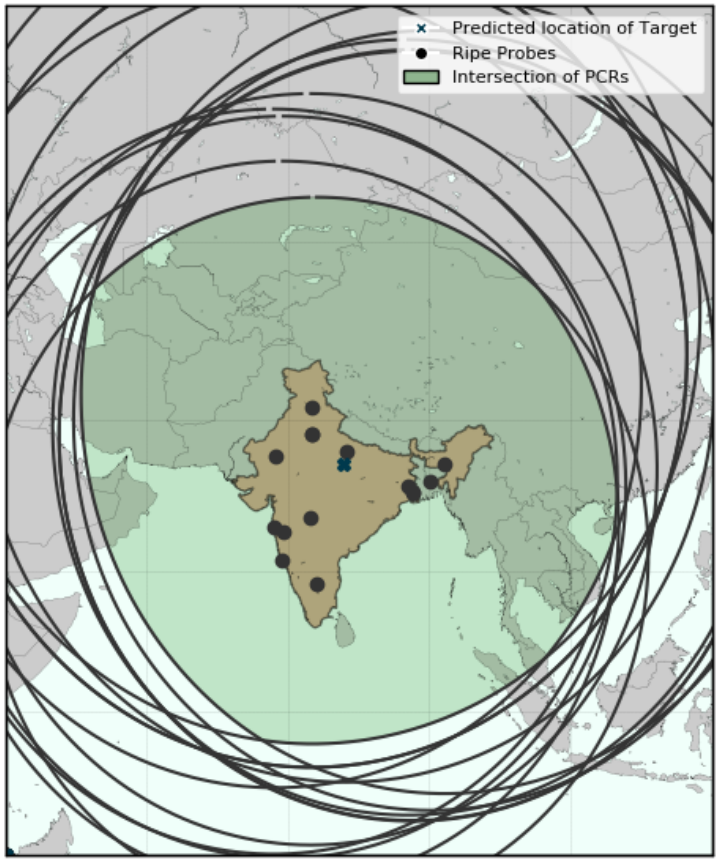

Figure 6.3: An IP address located outside India. The entire country (India) is well contained in the intersection area.

US), most of the PCRs might not intersect (as shown in Figure 6.5). But, since majority of these circles are within the nation's boundary, we ascribe this IP address as "inside".

6.3.2 Applying R-CBG for Different Countries

We selected 5 different nations namely IN, IR, SA, BR and US for our analysis. Our goal was to identify what fraction of Alexa top-'n' websites reside within these five nations by using R-CBG. For each country under consideration, we enumerate the steps taken for the same:

1. We selected Alexa top-1k websites, resolved them from a RIPE node located inside the country under consideration, and recorded their IP addresses.

2. Next, we individually selected the reference nodes, such that they were evenly distributed and located either inside, or close to, the geographic boundary of the nation.

Figure 6.4: Detecting IP anycasting. Dots represent the RIPE probes and circles represent the distance estimated to the IP address based on speed of light constraint.

3. The "bestline" (ref. Subsection 6.2.3) is computed using the RTT between the reference nodes (probes) and their (already known) geodesic distances. The probes `ping` each other and the corresponding RTTs are recorded.

4. Then the reference nodes `ping`-ed the targets, and the corresponding RTTs were also recorded.

5. We use the "bestline" and the recorded RTT (from the previous step 4) to estimate the distance between the reference nodes and the targets.

6. Using the reference nodes as centers and distances as radii (from the previous step 5), PCRs are drawn on world map. The intersection region of these PCRs and its centroid is computed.

7. We use the intersection region of PCRs, the centroid and the nation's boundary coordinates, along with the four-point heuristic, explained in Subsection 6.3.1, to decide if the target is positioned "inside" the nation or not.

Caveat: In cases where all PCRs do not intersect, we select the intersection region formed by maximally intersecting PCRs. There are two possible explanations. Firstly, a reference node experiencing heavy congestion may underestimate the distance to the target, leading to a smaller

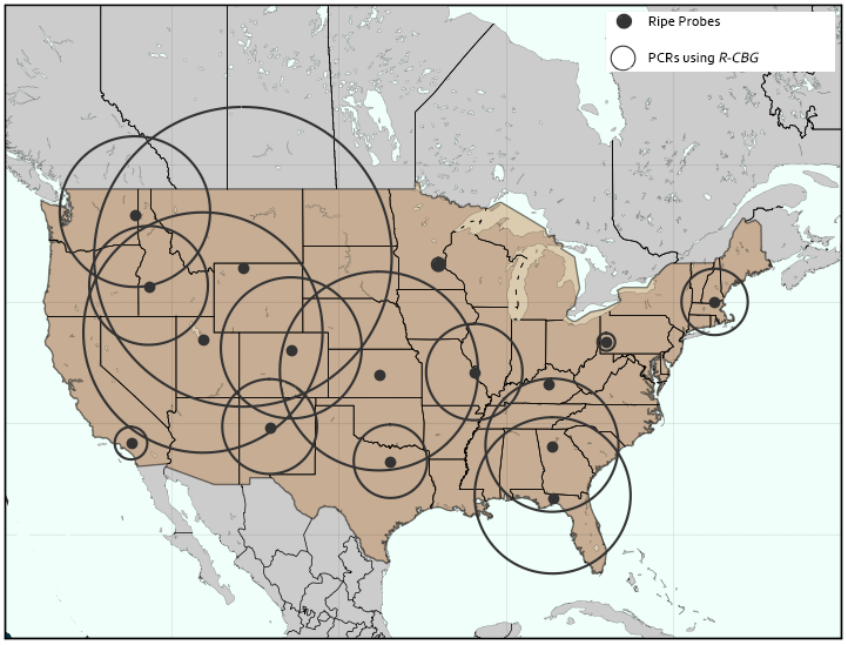

Figure 6.5: A single IP address is anycasted at multiple locations within US itself.

PCR [190]. Secondly, as already explained, this may be a case where the target IP is anycasted within the same nation. Maximally intersecting PCRs make R-CBG agnostic to such pitfalls.

6.3.3 Selection of Reference Nodes

We rely on RIPE Atlas for selecting our reference nodes. It offers two types of nodes — *anchors* and *probes*. Anchors are stable machines that regularly `ping` one another. On the contrary, probes are machines which may (or may not) be available all the time and might not respond to `ping` requests. Hence, we preferred selecting anchor nodes. However due to their scarcity in our tested countries, we also included a few stable probes.

While selecting probes we ensured the following: (i) Assuming target is within the nation, the probes are evenly distributed across, potentially surrounding the target. (ii) The probes are stable and respond to `pings`. To select the stable probes, we `ping`-ed them from our university machine for one week. We tested their connectivity 5 times a day, each time with 5 `ping` packets. Those probes which responded to more than 90% of the `ping` requests, qualified as stable reference nodes.

6.4 Data Collection and Results

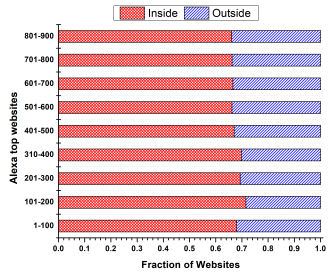

Figure 6.6: Fraction of websites located inside/outside Iran.

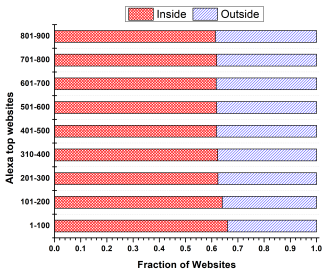

Figure 6.7: Fraction of websites located inside/outside India.

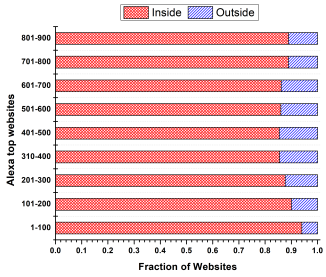

Figure 6.8: Fraction of websites located inside/outside Saudi Arabia.

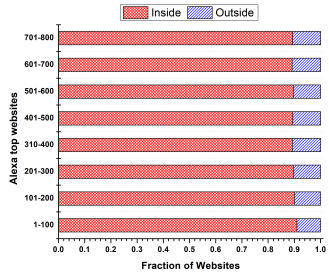

Figure 6.9: Fraction of websites located inside/outside Brazil.

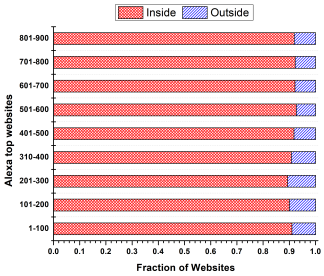

Figure 6.10: Fraction of websites located inside/outside United States.

Countries			Predicted		Accuracy (%)
			IN	OUT	
India (IN)	Actual	IN	19	2	89.73
		OUT	17	147	
Iran (IR)		IN	27	1	94.12
		OUT	5	69	
Saudi (SA)		IN	47	1	89.87
		OUT	7	24	
Brazil (BR)		IN	33	3	90.71
		OUT	10	94	
USA (US)		IN	617	4	92.41
		OUT	45	213	

Figure 6.11: Confusion Matrix for different countries.

6.4.1 Validating R-CBG

The aim of R-CBG is to determine if a target IP address resides in a country or not. To gauge its accuracy, we compared the outcomes of R-CBG against the domiciles of RIPE probes (known *a priori*). To that end, we first individually selected reference nodes in and around the nation under consideration. Further, for the targets, we considered all the RIPE nodes upto 5000 Km radius, from the country's approximate geographic center. To check their connectivity we `ping`-ed all of them. Those which responded were finally selected as targets, with their domiciles as the ground truths. This ensured that we have sufficient targets both inside and outside the nation.

On an average 91% of these targets, were correctly classified by R-CBG as inside (or outside) the chosen nation. *E.g.* corresponding to US, we identified total of 879 RIPE nodes (responsive to pings, and within 5000 Kms), that were selected as targets. Of these, 621 were hosted inside and 258 were outside. R-CBG correctly classified 617/621 nodes as inside, whereas 213/258

as outside. About 45 nodes which were actually hosted outside, were miss-classified as inside. Most of these were located near the border of US (either Mexico or Canada). Confusion matrix in figure 6.11 represents similar trends for all the five nations under consideration.

Caveat: It must be noted that before applying R-CBG to any new country, its accuracy needs to be tested again. For example, for very small countries like Switzerland, R-CBG might not result in high accuracy. The intersection region formed by the PCRs would be so large that it might always subsumes that entire country, irrespective of the location of target IP address (inside or outside the nation). This might violate the rationale behind our four-point heuristics.

6.4.2 Multilaterating Alexa Websites

Having established the accuracy of R-CBG, we used it to test if a websites is located inside the country or not. Figures 6.6, 6.7, 6.8, 6.9 and 6.10 show the fraction of websites that are located inside and outside the nation's boundary. *E.g.,* for Brazil, we observed that $\approx 89\%$ of Alexa top-1k websites are located inside its geographic periphery. Similarly, for Iran $\approx 66\%$, India $\approx 61\%$, Saudi Arabia $\approx 86\%$ and US $\approx 92\%$ of the websites were found to be located inside. However it must be noted that, in all the five countries, there were significant number of websites (atleast 8% of Alexa top-1k websites) that were certainly located *outside* the nation's boundary. Thus clients residing in these nations may use such sites with anti-censorship approaches like Decoy Routing.

Interestingly, we also observed that the sites hosted outside varied evenly in their popularity. *E.g.* in Saudi Arabia $\approx 8\%$ of the top-100 websites were found to be located outside. Similar trends were also observed for top 800-900 websites. Thus, selectively censoring the relatively less popular sites may not significantly impact anti-censorship systems that rely on websites positioned outside, however would hinder them.

6.4.3 Multilaterating Alexa Websites when Resolved from Outside the Nation

As already described, for majority of the websites, a client obtains a corresponding IP address in its own nation (likely due to CDNs). This might render anti-censorship approaches like Decoy Routing, Meek, CacheBrowser *etc.* ineffective. However, it can be argued that a client can still use these approaches, if it somehow obtains IP addresses corresponding to foreign front-ends of the same Alexa website. *E.g.,* if a client uses such IP addresses for Decoy Routing, the Decoy Routed packets may cross the nations network boundary, eventually being intercepted by the Decoy Router.

To obtain IP addresses outside of the censorious regime we individually resolved Alexa top-1k websites from a host (we control) in non-censorious country (Ireland) and recorded the IP addresses. Ideally these IP addresses should be located outside the censorious nation and may be used by the aforementioned anti-censorship approaches. Thus, to test our hypothesis, we multilaterated the resulting IP addresses from each of these countries, choosing the original set of reference nodes positioned within the said countries (as already mentioned in Subsection 6.3.3).

For each country, Figure 6.12 represents the number of country specific Alexa top-n websites, identified to be hosted inside, when domains were resolved from within the censorious nation, and from a non-censorious foreign nation. *Scenario A* is when websites were both resolved and multilaterated from the said country. Whereas, *Scenario B* corresponds to websites being resolved externally, but multilaterated using reference nodes inside the country.

Ideally, the websites which were earlier ascribed as inside the country (in Subsection 6.4.2), should have now been reported as outside. However, to our surprise, we observed no significant differences. The total number of websites hosted inside do not differ much in both the figures. *E.g.* in India 495/800 websites were inside in *Scenario A*, whereas in *Scenario B* it was 438/500. This implies that either majority of the websites were anycasted or non-CDN hosted (positioned only at a single location) within the bounds of the censor. This small difference (7.12%) is likely due to DNS based CDNs (explained in detail in the next subsection). Thus we now describe our approach to identify which type of CDN a website is using.

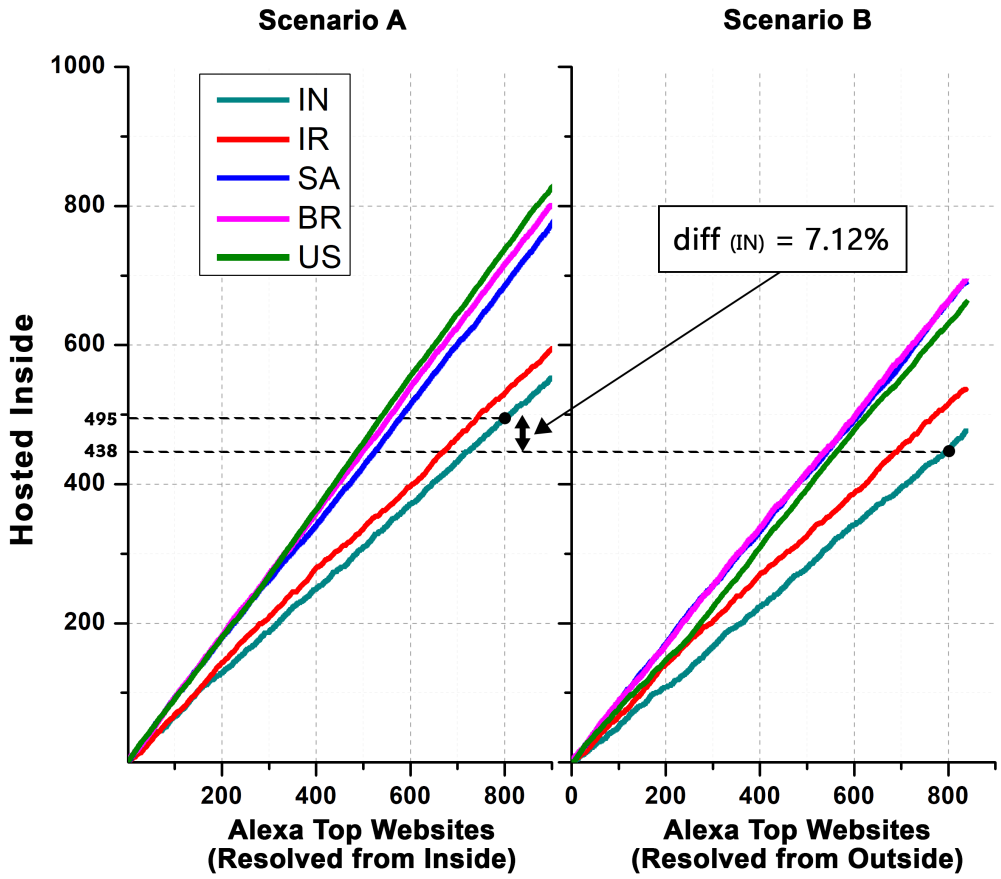

Figure 6.12: Number of websites located inside when resolved from within and outside the nation.

6.4.4 Identifying Type of CDN

Anycasted IP addresses are announced at multiple locations across the globe. Geolocating such IPs, using SOL based multilateration (ref. Subsection 6.2.3) would never yield an intersection of *all* PCRs. We use this observation to differentiate between anycast and other forms of hosting. This observation holds valid because probe packets of different reference hosts would be routed to their (likely) closest anycasting site. Thereafter, employing SOL constraint one estimates maximum possible distance travelled by a packet, in the observed RTT. Using these distances to multilaterate anycasted IP addresses would lead to zero (or very few) overlapping circles (ref. Figure 6.4).

To identify the different types of CDN hosting (for the Alexa-1k websites), we selected 25 globally distributed RIPE nodes and resolved each of the websites from them. A website that

resolved to the same IP address across all the probes, was multilaterated using the SOL principle. The presence of an intersection region of *all* PCRs, indicates that the website is positioned at only one location. The absence indicates anycast hosting.

On the other hand, a website that resolved to multiple IP addresses from the different probes, very likely uses DNS-based CDN. However, anycasting *also* allows a site to have multiple IP addresses which may simultaneously be advertised from various geographic locations (anycasting sites).

To differentiate the two, we randomly chose a single IP address for all such websites, and multilaterated it using SOL principle. Again, the presence of an intersection of *all* PCRs indicates that the IP is positioned only at one location, contraindicating anycasting, confirming DNS-based hosting. Figure 6.13 schematically describes the approach.

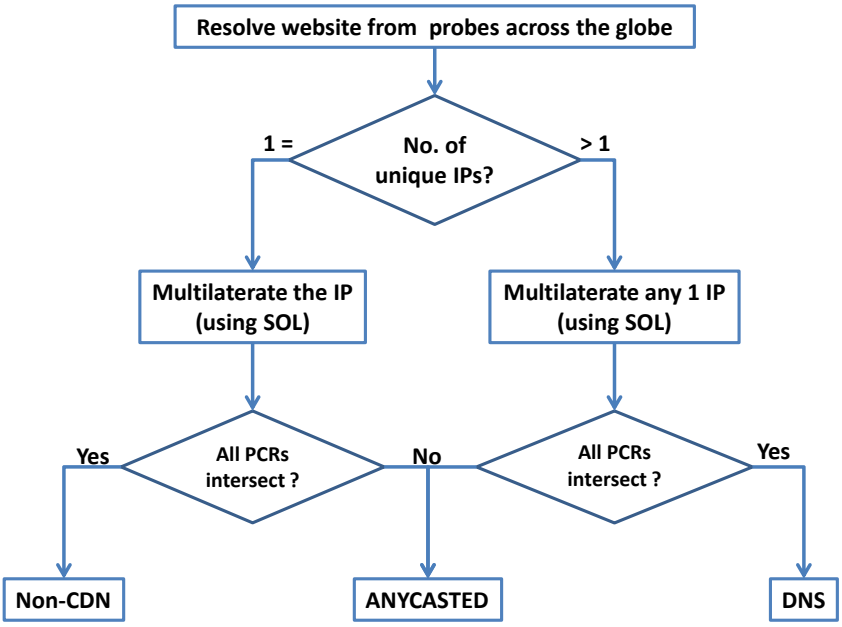

Figure 6.13: Identifying type of CDN used.

Table 6.1 represents the type of CDNs a website is using. Columns 1, 3 and 5 represent fraction of websites located inside the nation, whereas remaining represent websites hosted outside (coloured as green). Anti-censorship approaches (like Decoy Routing, CacheBrowser) require websites to be located outside the censor's boundary. However, websites using *DNS-based CDNs* which are located *inside* (column 1) can also be used for such purposes. These

websites very likely have front-ends spread across the globe. If a client (somehow) obtains an IP address of a foreign front-end, its request will cross the national boundary. Thus, websites using DNS-based CDNs located inside can also be considered as a viable option for such anti-censorship approaches.

Additionally, the Table 6.1 also explains the reason for minor differences (*e.g.* 7.12% for IN) reported in Figure 6.12. It corresponds to the fraction of websites using DNS based CDNs, hosted inside. When websites are resolved from a foreign host would likely map to their respective foreign front-end IP address, while others (anycast and non-CDN based) would not. For example, the IP addresses for 61.88% websites for India, when resolved internally, were identified to be inside. This number dropped to 54.8% when these sites were resolved externally. This difference is close to the fraction of websites using DNS based CDNs (*i.e.,* 8.12%), hosted inside.

Nations	CDN				Non-CDN	
	DNS		ANYCAST			
	Hosted In (%)	Hosted Out (%)	Hosted In (%)	Hosted Out (%)	Hosted In(%)	Hosted Out (%)
IN	8.12	6.77	35.92	11.81	19.68	17.71
IR	1.31	2.42	11.78	15.01	52.87	16.62
SA	8.41	1.19	57.44	3.66	20.80	8.51
BR	9.36	2.07	52.88	2.72	27.97	5.01
US	14.97	3.49	56.31	2.67	20.62	1.95

Table 6.1: Type of CDNs used by Alexa top-1k websites.

6.5 Inferences from Results

6.5.1 Nation State Hegemony

Our results reveal that for five countries under consideration, on an average $\approx 80\%$ of Alexa top-1k websites are located within the country. Moreover, this trend remains same when we extended our analysis to Alexa top-5k websites (ref. Subsection 6.6.1). This clearly indicates that popular web traffic from clients does not cross their geographic boundaries. This contradicts

the notion that certain powerful countries may be hegemonic over the Internet and might surveil the traffic of "underserved" nations transiting through them.

However, it can be argued that even if websites are located inside the nation, web requests may follow *tromboning paths* — paths that originate and end in the same country, but transit a foreign country [9]. Our research indicates that, even if tromboning paths exist, they are rare. A packet following a tromboning path (relatively longer path) would very likely result in higher RTT, in comparison to a non-tromboning path [219]. Such high RTTs would have resulted in wider PCRs, resulting in large intersection regions. In such cases, R-CBG would have incorrectly classified inside nodes as outside (ref. point 4 of Subsection 6.3.1). But, this is not the case as R-CBG has high accuracy ($> 90\%$). *E.g.* in US, R-CBG correctly identified 617 (out of 621) internal RIPE nodes. Such observations likely indicate the absence of tromboning paths.

Such inferences convey that the client's traffic is majorly localized in its own country, and rarely follows a hierarchical path [167]. Thus, we believe that CDNs have resulted in "Internet flattening", also reported by others [220, 180, 221].

6.5.2 Hindrances to Anti-Censorship

Decoy Routing: Decoy Routing assumes that there are some unblocked "overt" websites, located outside the censorious regimes. Thus, it is expected that when a client sends web requests (carrying special cryptographic patterns) to these "overt" websites, they cross the censors' nation boundary. *En route* these requests are intercepted by Decoy Routers (DRs), also positioned outside. Based on the signatures, the DR identifies the packets and diverts them towards the intended "covert" destination. Our results reveal that front-ends of $\approx 80\%$ of the popular websites are located within the client's own nation. This poses challenges for the placement of DRs [154, 51, 181].

Meek: As already mentioned, a Meek client engages in an apparent HTTPS communication with a non-censored site (the front-end). It actually bears a request to the blocked domain in the HTTP Host field. The front-end (inside the censor's boundary) is oblivious to censorship

and attempts to fetch filtered content from its back-end server. However, since front-ends are themselves located within the censor's control, their communication with back-end server may be intercepted by the adversary. Yet again, our research reveals that front-ends of $\approx 80\%$ of the popular websites are located within the client's nation. Additionally, a censor might coerce the CDN provider preventing the front-ends (positioned in its boundary) from responding to request to blocked domains.

CacheBrowser: When clients access CDN-based web content, the requests are generally directed to the closest front-ends (often located in clients' country). However, Zolfaghari *et al.* [209] reported that powerful censors coerce CDN providers to filter content on front-ends located in their bounds.

CacheBrowser (and its successor CDNReaper) aims to disrupt this dynamic. It involves direct communication with IP addresses of foreign front-ends, rather than relying on regular DNS resolutions. In general anycast based CDNs assign the same IP address to all their front-ends. Thus, for anycasted websites, it cannot be used; the request would never leave the clients' countries. *This implies that CacheBrowser can only be used to access websites that use DNS based CDNs.* Our results reveal that only a small fraction of Alexa top-1k websites rely on DNS based CDNs *i.e* $\approx 11\%$ (ref. Table 6.1). CacheBrowser can be used to access only these websites.

But, it can be argued that a censor might avoid blocking Alexa top-1k websites due to their popularity. Thus, we also tested if CacheBrowser can be employed for accessing potentially blocked websites (Citizen Lab's [222] country specific lists). It is evident from Figure 6.14, that majority of the such websites (for all five countries) were not using DNS-based CDNs, and thus may be inaccessible using CacheBrowser. *E.g.* in BR, out of 2769 potentially blocked URLs, CacheBrowser can only be used to access about 525 of them (*i.e* $\approx 18\%$).

CovertCast: It relies on popular live video streaming services (*e.g.* YouTube) to secretly transport the content of blocked websites to clients residing in censorious regimes. Popular streaming

video services like YouTube use CDNs and operate via front-ends (in censor's boundary). Thus, the adversary might coerce these services to stop support for CovertCast, particularly via front-ends positioned within its control. In our results, we observed popular live stream supported sites (*e.g.* YouTube, Facebook, Instagram) have front-ends hosted within the countries under test.

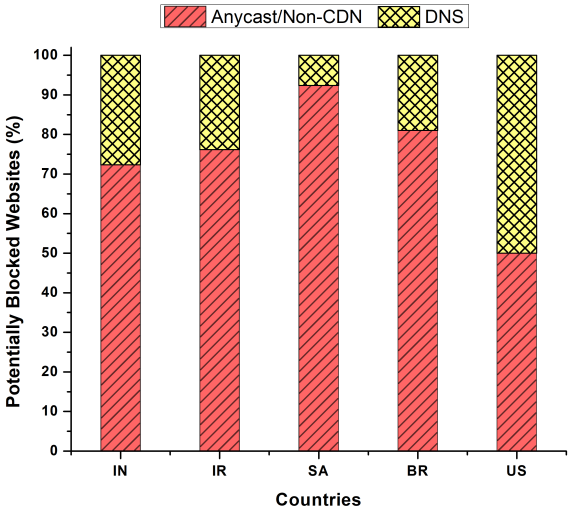

Figure 6.14: Type of CDNs used by potentially blocked websites hosted.

6.6 Discussion

6.6.1 Projection to Alexa top-5K Websites

In previous sections, we identified that majority of Alexa top-1k sites were located within the nations under test. However, one can question if our analysis holds good for other Alexa websites as well. Thus, we further attempt to identify the locations of Alexa top-5k sites w.r.t the individual countries.

The ideal case would have been to use R-CBG to multilaterate each of these websites. However, the multilateration process is *costly* in terms of RIPE credits and time consumption. In total we required multilaterating $\approx 25K$ websites. In the absence of sufficient credits, we abandoned this idea and relied on RTT[1] as a gross-metric to identify the location of these websites. We now

[1]Measuring RTT using `ping` reduces the credit requirement by ≈ 20 times in comparison to running R-CBG for a single target.

explain, why RTT can be used as a metric for the same.

Websites locality using RTT profiling: As already described in Subsection 6.2.3, RTT by and large correlates to distance. For most of the Alexa top-1k sites, the ones hosted inside have relatively smaller RTTs than those hosted outside. *E.g.* Figure 6.15 shows the distribution of these RTTs, recorded using a single RIPE probe in Iran. Evident from the figure, there are *two* distinguishable categories corresponding to websites located inside and outside. Similar trends were also observed for other countries (ref. Appendix 7).

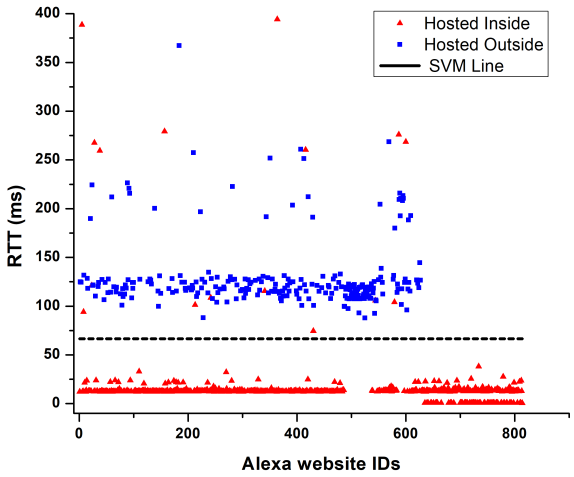

Figure 6.15: RTT scatter plot for a probe in Iran.

From such observations, we believe that RTT *alone* can be used to determine whether websites are internal (or external) to the nation. To correctly disambiguate the two categories, we used Support Vector Machine (SVM) classifier [223]. To SVM classifier we provide a set of labelled data points. The set contains Alexa top-1k website labelled as inside or outside (outcomes of R-CBG in Subsection 6.4.2) and their corresponding RTT values obtained from the reference node. SVM then creates the best possible line that divides the set of RTT into two distinct categories — inside and outside. To test the effectiveness of SVM we used K-fold cross-validation for a reference node (K=5) [224]. It was then repeated for all reference nodes in a country.

K-fold cross validation (K-5), divides the samples into five equal sized partitions. Of the five

partitions, four are used for training, the remaining one is used for testing the classifier. The entire process is repeated five times, selecting a different (non-repetitive) testing partition in every iteration. Then, the average classification accuracy (and corresponding standard deviation) of the five iterations is computed. For these nodes the cross-validation process observed an accuracy of $85 - 99\%$, barring a few outliers[2].

Next, we computed the Coefficient of Variation (CoV) for all the reference nodes. The reference node with least CoV was considered as the most consistent one. Finally, for every country we selected the most consistently accurate reference node for further analyzing the locality of Alexa top-5k websites. Figure 6.15 shows the SVM line (threshold) for such a reference node in Iran. Similar thresholds were also established for other nations (ref. Appendix .3).

Fraction of Alexa top-5k websites hosted inside (or outside) the country: We measured the RTT for Alexa top-5k sites from the reference nodes with highest consistency Based on the previously trained SVM classifiers, we predicted what fraction of 5k websites reside inside or outside for each country (ref. Fig. 6.16). By and large the fractions of websites hosted inside remain the same for both Alexa top-5k and top-1k. This yet again dispels the notion of nation state hegemony even for a larger set of websites. Additionally, this shows that most of these sites cannot be used by anti-censorship systems like Decoy Routing.

Interestingly for Iran, a large fraction of the more popular websites (ranked below 1k) were hosted inside, while the less popular ones were hosted outside (ref. Fig. 6.16). This anomalous behaviour can be explained as follows. From Table 6.1, we observe that majority (52.87%) of the Alexa top-1k websites were located inside and used non-CDN infrastructure. On the other hand, only a small fraction ($\approx 13\%$) of popular websites that were found inside, used CDNs. This indicates, low CDN presence inside Iran. As a consequence, for less popular websites (ranked above 1k) running on CDNs, the lack of internal front-ends may force requests to exit the nation boundary. This is likely why we observe an inversion in the trend.

[2]For each country, for around 2-3 reference nodes, we observed lower accuracy ($\approx 50 - 60\%$). This may be attributed to variable congestion experienced by these nodes, resulting in dramatic RTT variations.

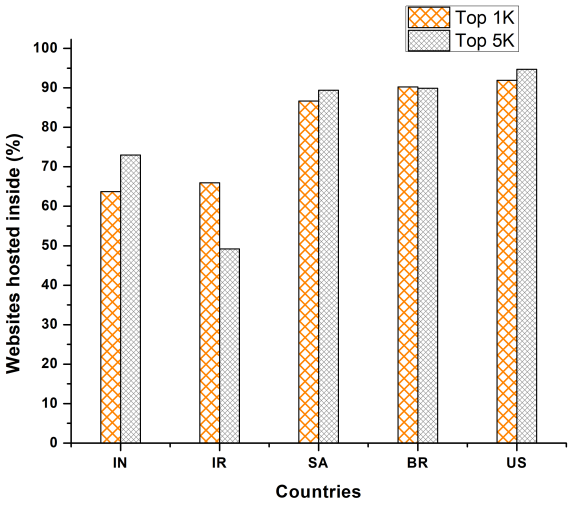

Figure 6.16: Percentage of Alexa top-5k websites hosted inside the country.

6.6.2 Decoy Routing via Parallel (Leaf) Web Connections

For all the countries we observed that a majority of the web request (to Alexa websites) never crossed the nation boundaries. As already discussed, this may hinder the functioning of anti-censorship systems like Decoy Routers (ref. Section 6.5.2).

Response for typical web requests bear HTML code embedded with URLs for content like CSS and images (required to render the pages). Web browsers establish individual (parallel) connections corresponding to each of these URLs. *Some of these parallel connections may be utilized for Decoy Routing*, if they terminate at IP addresses located outside the censors' control.

Most parallel connections do not cross the nation boundary. For example, when we analyzed the web transactions corresponding to Alexa top-100 websites for India, we found that for 23 websites, the parallel connections terminated outside, even when the original web requests did not. We observed that most of these websites resulted in very few parallel connection terminating outside the nation (mode 1, median 3). However, for one particular website, this value was as high as 32. In future, Decoy Routing systems can evolve to make use of such parallel connections.

6.6.3 Comparison with Popular Geolocation Databases

It is known that popular geolocation databases are prone to errors at city level. However, it is believed that at country level they are relatively error free [225]. *E.g.*, very recently Edmundson *et al.* [9], used Maxmind for IP-to-country mapping. Using this they reported that majority of the Internet paths (to Alexa top-1k websites) crossed the national boundary. We thus compared the accuracy of Maxmind GeoIP2 database [211] with R-CBG. To do so, we compared the country level information for Alexa top-1k sites (for each of the countries) derived from the database, against the results obtained by multilaterating with R-CBG. Our results show (ref. Figure 6.17) that geolocating IP addresses of popular websites using Maxmind results in large errors. For instance, in SA, Maxmind report $< 15\%$ of the websites to be located inside the country itself, whereas our algorithm ascribed 90% of them to be inside. This is because, often these databases rely on static annotations [212], which map an IP address to the country where parent AS is headquartered. Thus, relying on such information [9], to determine the country through which traffic transits, seems inaccurate. Hence the notions that "powerful" countries intercept large fraction of traffic originating from underserved nations, is unfounded.

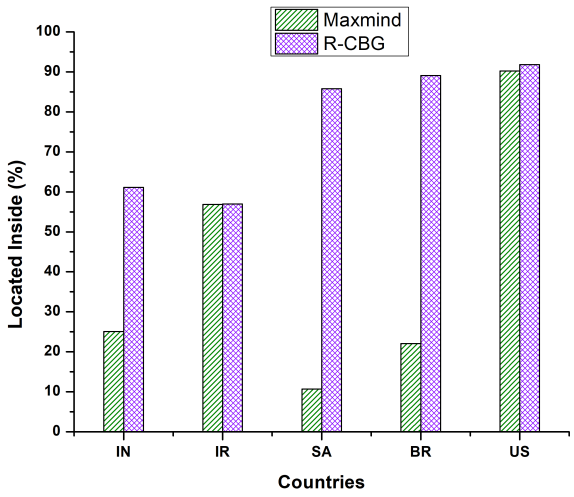

Figure 6.17: Comparison of Maxmind with R-CBG.

6.6.4 Selection of Countries

We studied few nations which are known to censor and surveil network traffic [11]. Further, R-CBG requires about 15 probes in a geographic vicinity of a nation under study. Unfortunately, the RIPE nodes are concentrated in EU and North America [110, 190]. We thus chose nations where atleast 15 stable nodes were available. Moreover, we restricted our study to five nations as our objective was to judiciously use the limited RIPE "credits"[3].

We believe that these nations represent a diverse set in terms of geo-political power, Internet infrastructure and free and open communication.

6.6.5 Selection of Reference Nodes

The accuracy of R-CBG depends on the number of reference nodes and their geographic proximity to the target. We chose the reference nodes within (or close) to a country. Thus, selecting enough nodes ensured that the targets were close to a considerable fraction. We empirically observed that atleast 15 nodes were required in each of the countries for R-CBG to correctly multilaterate the target.

Previous authors [189] however reported that about 30 reference nodes were required to geo-locate (using CBG) the targets. But, they chose the reference nodes at the continent level while we chose within the country (closer to the target).

6.6.6 Selection of Target Websites for R-CBG

We selected country specific Alexa top-n websites (n=100, 200...1000) for two reasons. Firstly, several anti-censorship approaches like Decoy Routing rely on these sites and require them to be positioned beyond the censors' control. Our analysis aids assessing the feasibility of using such systems in the nations considered. Secondly, these websites are a representation of the actual web traffic of users. Others such as Cisco *Umbrella* [227] and *Majestic Million* [228] are derived from indirect sources like DNS queries and URLs embedded in website ads, often rendered

[3]The details on how RIPE credit system works can be found at [226]

or accessed without the users' control. Our choices are in accordance to recommendations by Scheite *et al.* [47].

6.6.7 Stability of Our Results

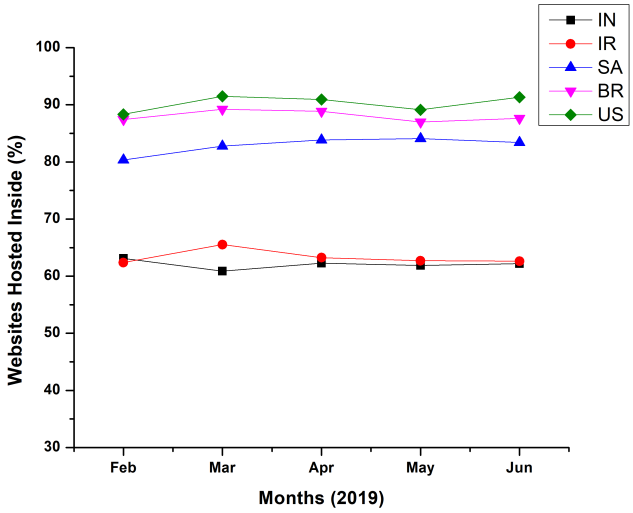

Figure 6.18: Location stability for Alexa top-1k websites.

Scheite *et al.* [47] report that list of popular websites vary over time. Our measurements may be subject to such variations. Thus, we repeated our experiments for five consecutive months by selecting monthly snapshots of Alexa top-1k websites. As evident from Figure 6.18), the fraction of websites hosted inside roughly remained same over the said period. On an average, the ranks of about 812 websites remained less than 1000, across the five monthly snapshots.

6.6.8 Limitations

To identify whether a website resides within (or) outside the nation, R-CBG requires reference nodes to be placed evenly within a nation. To do so, we manually selected the reference nodes. To remove manual intervention, we plan to automate this process in future.

Additionally, our four-point heuristics which classify a given target as inside (or outside) the nation, depends on size of size of the country as well as the the intersection region. For a target placed inside a sufficiently large country (*e.g.* US), the intersection region produced by PCRs

(using evenly distributed nodes) very likely lie inside the country. However, for a relatively small country (like Switzerland), even a target residing inside, might be classified as outside. This might happen because intersection region might be so large that it subsumes the entire nation. Thus, for five countries under considerations, we first tested R-CBG against the ground truth and obtained atleast 90% accuracy for each of them.

6.7 Conclusion

The proliferation of CDNs on the Internet, have brought web content closer to the end-user. On the positive side, it has improved users' web experience. Additionally, this likely contradicts notions like nation state hegemony, *i.e.,* "powerful" nations intercepting traffic from "underserved" nations. We empirically confirm this contradiction.

On the negative side, this content "closeness" may enhance nation states' ability to coerce content providers for regulating access within their own boundaries. From our tests repeated for five consecutive months, we identified that vast majority (>80%) of popular web content (Alexa top-1k websites) is located within the nation states. To identify if a host is inside or outside a nation, we re-engineered Constraint Based Geolocation (CBG), a multilateration technique. In $\approx 91\%$ cases it correctly identified if a host is inside (or outside) a country. We call it Region Specific CBG (R-CBG).

The natural solution to evade surveillance (or censorship), involves proxy based systems. Sadly, these often bear easily identifiable traffic signatures (*e.g.* IP address). Newer alternatives, like Decoy Routing, CacheBrowser *etc.*, require access to popular website hosted outside the censors' control. Unfortunately, CDNs may hinder easy access to such websites. Our heuristics classified overwhelming majority of the Alexa top-1k websites as hosted on CDNs, within the clients' domicile. Interestingly, this trend persists for Alexa top-5k websites, when tested using a novel RTT based heuristic. Thus, neither conventional (proxy based), nor heterodox (relying on web traffic) approaches alleviate the predicament.

However, a small, yet significant set of websites ($\approx 20\%$), are hosted outside the censors'

boundaries and may be used by such systems. On the contrary, traditional proxy based systems, do not rely on CDNs. But the arms race to camouflage easily identifiable patterns [229], and efforts to discover them [203], continues.

Chapter 7

Conclusion

Surveillance is the business model of the Internet.

Bruce Schneier

With our research efforts summarized in this thesis we experimentally demonstrate that Internet cartographic techniques can aid censorship studies. More specifically, we demonstrate that inferences drawn from maps can be used to aid both censorship and its circumvention. For example, in India we require a cooperation of handful ASes (< 10) for achieving effective censorship. On the other hand, we require co-operation of 30 ASes to build a global infrastructure for newer anti-censorship approaches like Decoy Routing.

ASes are not a monolithic entity — they constitute hundreds to thousands of routers managed by an organization. We constructed router level maps of various ASes to identify what are the "choke points" in the AS itself *i.e.,* number of routers that have the potential to capture large fraction of the AS's traffic. This vital information can both be used (by the civil liberties organizations) and abused (by the censor). For example, we identified that around $11,000$ routers must be replaced with DRs (within the key 30 ASes) for a global Decoy Routing infrastructure — a worthwhile effort for the free Internet. Whereas, the same maps can be used by the censor to deploy middleboxes to achieve coherent and effective censorship. For instance, our research efforts reveal that India requires around $5,000$ routers for IP or DNS filtering.

Further, to analyze where the the middleboxes are placed within the ISPs we proposed two metrics *viz.,* Coverage and Consistency. Coverage implied fraction of network entities (within an ISP) that are involved in some form of censorship and Consistency implies the fraction of censorship infrastructure that is blocking the same content. For example, in MTNL we found that $> 85\%$ of total DNS resolvers hosted within the ISP were poisoned *i.e.,* they were engaged in DNS censorship (Coverage) and a single website was blocked by $< 50\%$ of the poisoned resolvers (Consistency).

Internet maps constructed *solely* from publicly available BGP data might lead to incorrect inferences. For example, BGP path(s) originating from a Iranian ISP, and terminating at `google.com` (headquartered in US) might traverse multiple international ASes. However, `google.com` is a managed using its own CDN, that also has presence in India. In practice, our research shows request from Iranian ISP's customers to `google.com` do not cross Iran's boundaries. Further we also confirm that 80% of Alexa top-1k websites are located within the nation states due to the proliferation of CDNs. Thus, inadvertently CDNs can provide more control to the censors as the web content is now localized within their own jurisdictions. Moreover, it poses additional operational challenges for newer anti-censorship approaches (like Decoy Routing) that rely on web content being hosted outside the censor's control.

Thus in future we aim to develop new Internet cartographic techniques that would also incorporate the impact of routing changes in Internet due to CDNs. This would provide more realistic topologies of the Internet — *i.e.,* we would be able to predict an actual path a packet would traverse before reaching its destination. For instance, we may augment (or improve) existing BGP paths based maps with traceroute paths, followed by geolocating the intermediate routers and the end destination. Such maps could be further used by anti-censorship community to study where censorship infrastructure is placed and how to route around it.

References

[1] L. Subramanian, S. Agarwal, J. Rexford, and R. H. Katz, "Characterizing the Internet hierarchy from multiple vantage points," in *INFOCOM 2002. Twenty-First Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, vol. 2. IEEE, 2002, pp. 618–627.

[2] U. N. G. Assembly, "Thirty-second session, human rights council, item 3. the promotion, protection and enjoyment of human rights on the Internet." [Online]. Available: https://www.ohchr.org/EN/HRBodies/HRC/RegularSessions/Session32/Pages/ResDecStat.aspx

[3] Article19, "Significant resolution reaffirming human rights online adopted." [Online]. Available: https://www.article19.org/resources.php/resource/38429

[4] "India proposes Chinese-style Internet censorship." [Online]. Available: https://www.nytimes.com/2019/02/14/technology/india-{I}nternet-censorship.html

[5] "India: Historic Supreme Court ruling upholds online freedom of expression." [Online]. Available: https://www.amnesty.org/en/latest/news/2015/03/india-supreme-court-upholds-online-freedom-of-expression/

[6] E. Carisimo, C. Selmo, J. I. Alvarez-Hamelin, and A. Dhamdhere, "Studying the evolution of content providers in the Internet core," in *2018 Network Traffic Measurement and Analysis Conference (TMA)*. IEEE, 2018, pp. 1–8.

[7] L. Gao, "On Inferring Autonomous System Relationships in the Internet," *IEEE/ACM Trans. Netw.*, vol. 9, no. 6, pp. 733–745, Dec. 2001.

[8] A. Houmansadr, E. L. Wong, and V. Shmatikov, "No direction home: The true cost of routing around decoys." in *NDSS*, 2014.

[9] A. Edmundson, R. Ensafi, N. Feamster, and J. Rexford, "Nation-State Hegemony in Internet Routing," in *Proceedings of the 1st ACM SIGCAS Conference on Computing and Sustainable Societies*, ser. COMPASS '18. New York, NY, USA: ACM, 2018, pp. 17:1–17:11. [Online]. Available: http://doi.acm.org/10.1145/3209811.3211887

[10] C. Bocovich and I. Goldberg, "Slitheen: Perfectly imitated decoy routing through traffic replacement," in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2016, pp. 1702–1714.

[11] "Freedom House," https://freedomhouse.org/.

[12] "Internet live statistics." [Online]. Available: http://www.{I}nternetlivestats.com/

[13] P. Levis, "The collateral damage of Internet censorship by DNS injection," *ACM SIGCOMM CCR*, vol. 42, no. 3, 2012.

[14] L. Subramanian, S. Agarwal, J. Rexford, and R. H. Katz, "Characterizing the Internet hierarchy from multiple vantage points," in *INFOCOM 2002. Twenty-First Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, vol. 2. IEEE, 2002, pp. 618–627.

[15] G. C. Feng and S. Z. Guo, "Tracing the route of China's Internet censorship: An empirical study," *Telematics and Informatics*, vol. 30, no. 4, pp. 335–345, 2013.

[16] B. Liang and H. Lu, "Internet development, censorship, and cyber crimes in China," *Journal of Contemporary Criminal Justice*, vol. 26, no. 1, pp. 103–120, 2010.

[17] X. Xu, Z. M. Mao, and J. A. Halderman, "Internet censorship in China: Where does the filtering occur?" in *International Conference on Passive and Active Network Measurement*. Springer, 2011, pp. 133–142.

[18] B. Warf, "Geographies of global Internet censorship," *GeoJournal*, vol. 76, no. 1, pp. 1–23, 2011.

[19] B. Wagner, "Push-button-autocracy in Tunisia: Analysing the role of Internet infrastructure, institutions and international markets in creating a tunisian censorship regime," *Telecommunications Policy*, vol. 36, no. 6, pp. 484–492, 2012.

[20] J. Zittrain and B. Edelman, "Internet filtering in China," *IEEE Internet Computing*, vol. 7, no. 2, pp. 70–77, 2003.

[21] J. Qiu and L. Gao, "As path inference by exploiting known as paths," in *Global Telecommunications Conference, 2006. GLOBECOM'06. IEEE.* IEEE, 2006, pp. 1–5.

[22] D. Meyer, "The University of Oregon Routeviews Project." [Online]. Available: www.routeviews.org

[23] Anonymous, "The collateral damage of Internet censorship by DNS injection," *SIGCOMM Comput. Commun. Rev.*, vol. 42, no. 3, pp. 21–27.

[24] J. Zittrain and B. Edelman, "Internet filtering in China," *IEEE Internet Computing*, vol. 7, no. 2, pp. 70–77, 2003.

[25] R. Deibert, J. Palfrey, R. Rohozinski, J. Zittrain, and J. G. Stein, *Access denied: The practice and policy of global Internet filtering*. Mit Press, 2008.

[26] S. Wolfgarten, "Investigating large-scale Internet content filtering," 2006. [Online]. Available: shorturl.at/wzCZ0

[27] M. Dornseif, "Government mandated blocking of foreign web content," *arXiv preprint cs/0404005*, 2004.

[28] J. Wright, "Regional variation in Chinese Internet filtering," *Information, Communication & Society*, vol. 17, no. 1, pp. 121–141, 2014.

[29] S. Aryan, H. Aryan, and J. A. Halderman, "Internet Censorship in Iran: A First Look." in *FOCI*, 2013.

[30] Z. Nabi, "The anatomy of web censorship in pakistan." in *FOCI*, 2013.

[31] J.-P. Verkamp and M. Gupta, "Inferring mechanics of web censorship around the world." in *FOCI*, 2012.

[32] J. R. Crandall, D. Zinn, M. Byrd, E. T. Barr, and R. East, "Conceptdoppler: a weather tracker for Internet censorship." in *ACM Conference on Computer and Communications Security*, 2007, pp. 352–365.

[33] J. Zittrain, R. Budish, and R. Faris, "Herdict: Help spot web blockages," 2014.

[34] A. Sfakianakis, E. Athanasopoulos, and S. Ioannidis, "Censmon: A web censorship monitor," in *USENIX Workshop on Free and Open Communication on the Internet (FOCI)*, 2011.

[35] S. Burnett and N. Feamster, "Encore: Lightweight measurement of web censorship with cross-origin requests," *ACM SIGCOMM Computer Communication Review*, vol. 45, no. 4, pp. 653–667, 2015.

[36] R. Clayton, S. J. Murdoch, and R. N. Watson, "Ignoring the Great Firewall of China," in *International Workshop on Privacy Enhancing Technologies*. Springer, 2006, pp. 20–35.

[37] J. C. Park and J. R. Crandall, "Empirical study of a national-scale distributed intrusion detection system: Backbone-level filtering of html responses in China," in *Distributed Computing Systems (ICDCS), 2010 IEEE 30th International Conference on*, 2010, pp. 315–326.

[38] P. Winter and S. Lindskog, "How the Great Firewall of China is Blocking Tor." in *FOCI*, 2012.

[39] R. Govindan and H. Tangmunarunkit, "Heuristics for Internet map discovery," in *INFOCOM 2000. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, vol. 3. IEEE, 2000, pp. 1371–1380.

[40] H. Chang, S. Jamin, and W. Willinger, "Inferring AS-level Internet topology from router-level path traces," in *ITCom 2001: International Symposium on the Convergence of IT and Communications*. International Society for Optics and Photonics, 2001, pp. 196–207.

[41] Y. Shavitt and E. Shir, "Dimes: Let the Internet measure itself," *ACM SIGCOMM Computer Communication Review*, vol. 35, no. 5, pp. 71–74, 2005.

[42] K. C. Claffy, "The caida archipelago project." [Online]. Available: http://www.caida.org/projects/ark/

[43] H. V. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. Anderson, A. Krishnamurthy, and A. Venkataramani, "iplane: An information plane for distributed services," in *Proceedings of the 7th symposium on Operating systems design and implementation*. USENIX Association, 2006, pp. 367–380.

[44] K. Claffy, Y. Hyun, K. Keys, M. Fomenkov, and D. Krioukov, "Internet mapping: from art to science," in *Conference For Homeland Security, 2009. CATCH'09. Cybersecurity Applications & Technology*. IEEE, 2009, pp. 205–211.

[45] M. Luckie, B. Huffaker, A. Dhamdhere, V. Giotsas, and K. C. Claffy, "As relationships, customer cones, and validation," in *Proceedings of the 2013 Conference on Internet Measurement Conference*, pp. 243–256.

[46] "100 websites that rule the Internet," https://www.vodien.com/da-100-websites-rule-{I}nternet.php.

[47] Q. Scheitle, O. Hohlfeld, J. Gamba, J. Jelten, T. Zimmermann, S. D. Strowes, and N. Vallina-Rodriguez, "A long way to the top: significance, structure, and stability of Internet top lists," in *Proceedings of the Internet Measurement Conference 2018*. ACM, 2018, pp. 478–493.

[48] "AS relationships," http://www.caida.org/data/as-relationships/.

[49] "AS relationships," http://www.caida.org/data/as-relationships/.

[50] "What is cone size?" [Online]. Available: http://www.caida.org/data/as-relationships

[51] A. Houmansadr, E. L. Wong, and V. Shmatikov, "No direction home: The true cost of routing around decoys." in *NDSS*, 2014.

[52] J. Cesareo, J. Karlin, J. Rexford, and M. Schapira, "Optimizing the placement of implicit proxies," https://www.cs.princeton.edu/~jrex/papers/decoy-routing.pdf, 2012.

[53] E. Gregori, A. Improta, and L. Sani, "Isolario: a do-ut-des approach to improve the appeal of BGP route collecting," *arXiv preprint arXiv:1611.06904*, 2016.

[54] E. Gregori, A. Improta, L. Lenzini, L. Rossi, and L. Sani, "A novel methodology to address the Internet as-level data incompleteness," *IEEE/ACM Transactions on Networking (TON)*, vol. 23, no. 4, pp. 1314–1327, 2015.

[55] V. Giotsas, "Improving the accuracy of the internet cartography," Ph.D. dissertation, UCL (University College London);, 2014.

[56] D. Gosain, A. Agrawal, S. Sekhawat, H. B. Acharya, and S. Chakravarty, "Mending Wall: On the Implementation of Censorship in India," in *Proceedings of the 13th EAI International Conference on Security and Privacy in Communication Networks (SECURECOMM)*, 2017.

[57] J. R. Crandall, D. Zinn, M. Byrd, E. T. Barr, and R. East, "Conceptdoppler: a weather tracker for Internet censorship." in *ACM Conference on Computer and Communications Security*, 2007, pp. 352–365.

[58] R. MacKinnon, "Flatter world and thicker walls? blogs, censorship and civic discourse in China," *Public Choice*, vol. 134, no. 1-2, pp. 31–46, 2008.

[59] S. Guo and G. Feng, "Understanding support for Internet censorship in China: An elaboration of the theory of reasoned action," *Journal of Chinese Political Science*, vol. 17, no. 1, pp. 33–52, 2012.

[60] S. Aryan, H. Aryan, and J. A. Halderman, "Internet censorship in Iran: A first look," in *Presented as part of the 3rd USENIX Workshop on Free and Open Communications on the Internet*. Berkeley, CA: USENIX, 2013. [Online]. Available: https://www.usenix.org/conference/foci13/{I}nternet-censorship-iran-first-look

[61] Z. Nabi, "The Anatomy of Web Censorship in Pakistan," in *Presented as part of the 3rd USENIX Workshop on Free and Open Communications on the Internet*. Berkeley, CA: USENIX, 2013. [Online]. Available: https://www.usenix.org/conference/foci13/workshop-program/presentation/Nabi

[62] "Inferring mechanics of web censorship around the world," in *Presented as part of the 2nd USENIX Workshop on Free and Open Communications on the Internet*. Berkeley, CA: USENIX, 2012.

[63] "Censorship in Sweden," https://www.dangerandplay.com/2016/01/29/sweden-caught-censoring-the-{I}nternet-1984-style/.

[64] "Censorship in France," http://www.laquadrature.net/en/french-parliament-approves-net-censorship.

[65] "Number of Indian Internet users," http://www.{I}nternetlivestats.com/{I}nternet-users-by-country/.

[66] "India is partly free by Freedom House," https://freedomhouse.org/report/freedom-net/2011/india.

[67] "Porn websites blocked in India: Government plans ombudsman for online content," http://gadgets.ndtv.com/{I}nternet/news/porn-websites-blocked-in-india-government\-plans-ombudsman-for-online-content-723485.

[68] "Right to information, a citizen gateway," http://rti.gov.in/.

[69] J. Dalek, B. Haselton, H. Noman, A. Senft, M. Crete-Nishihata, P. Gill, and R. J. Deibert, "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, 2013, pp. 23–30.

[70] B. Jones, N. Feamster, V. Paxson, N. Weaver, and M. Allman, "Detecting DNS root manipulation," in *International Conference on Passive and Active Network Measurement*. Springer, 2016, pp. 276–288.

[71] H. Ballani, P. Francis, and X. Zhang, "A study of prefix hijacking and interception in the Internet," *SIGCOMM Comput. Commun. Rev.*, vol. 37, no. 4, pp. 265–276, Aug. 2007. [Online]. Available: http://doi.acm.org/10.1145/1282427.1282411

[72] Q. Jacquemart, "Towards uncovering BGP hijacking attacks," Ph.D. dissertation, Télécom ParisTech, 2015.

[73] R. Mahajan, D. Wetherall, and T. Anderson, "Understanding BGP misconfiguration," in *ACM SIGCOMM Computer Communication Review*, vol. 32, no. 4. ACM, 2002, pp. 3–16.

[74] F. Le, S. Lee, T. Wong, H. S. Kim, and D. Newcomb, "Detecting network-wide and router-specific misconfigurations through data mining," *IEEE/ACM transactions on networking*, vol. 17, no. 1, pp. 66–79, 2009.

[75] "Censorship in venezuela: Over 370 Internet addresses blocked." [Online]. Available: https://panampost.com/pedro-garcia/2016/07/20/censorship-in-venezuela-over-370-{I}nternet-addresses-blocked/

[76] C. Stevenson, "Breaching the Great Firewall: China's Internet censorship and the quest for freedom of expression in a connected world," *BC Int'l & Comp. L. Rev.*, vol. 30, p. 531, 2007.

[77] "Service Providers List – Telecom Regulatory Authority of India," http://www.trai.gov.in/Content/ProviderListDisp/3_ProviderListDisp.aspx.

[78] "Herdict:Help Spot Web Blockages," http://herdict.org/.

[79] N. Spring, R. Mahajan, D. Wetherall, and T. Anderson, "Measuring isp topologies with rocketfuel," *IEEE/ACM Trans. Netw.*, vol. 12, no. 1, pp. 2–16, Feb. 2004.

[80] R. Dingledine, N. Mathewson, and P. Syverson, "Tor: The second-generation onion router," DTIC Document, Tech. Rep., 2004.

[81] "The Internet censorship saga in India," https://{I}nternetdemocracy.in/2012/03/the-{I}nternet-censorship-saga-in-india/.

[82] "Websites blocked by Indian government," http://sflc.in/wp-content/uploads/2015/12/censorship.-2012-2015.pdf.

[83] "India's Supreme Court strikes down controversial Internet censorship law," https://techcrunch.com/2015/03/23/indias-supreme-court-strikes-down-controversial-{I}nternet-censorship-law/.

[84] "Censorship is India by India times," http://telecom.economictimes.indiatimes.com/tele-talk/{I}nternet-censorship-regulating-india-s-{I}nternet/1369.

[85] "Opennet initiative," https://opennet.net/.

[86] "ONI report for India," https://opennet.net/research/profiles/india.

[87] "Court cases regarding Internet censorship," https://opennet.net/news/india-court-summons-google-facebook-microsoft-executives.

[88] "Govt of India wants 32 URLs, including Dailymotion, Vimeo and Github, banned," http://indianexpress.com/article/technology/social/government-wants-32-urls-including-dailymotion-vimeo-banned-in-india/.

[89] "830 more websites blocked in India, many torrent links in list," http://indiatoday.intoday.in/technology/story/830-more-websites-blocked-in-india-many-torrent-links-in-list/1/748565.html.

[90] H. Haddadi, M. Rio, G. Iannaccone, A. Moore, and R. Mortier, "Network topologies: inference, modeling, and generation," *IEEE Communications Surveys Tutorials*, vol. 10, no. 2, pp. 48–69, Second 2008.

[91] "University of Oregon Route Views Project," http://www.routeviews.org/, 2000.

[92] L. Gao, "On Inferring Autonomous System Relationships in the Internet," *IEEE/ACM Trans. Netw.*, vol. 9, no. 6, pp. 733–745, Dec. 2001. [Online]. Available: http://dx.doi.org/10.1109/90.974527

[93] H. V. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. E. Anderson, A. Krishnamurthy, and A. Venkataramani, "iplane: An information plane for distributed services," in *Proceedings of* 7^{th} *USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, November 2006, pp. 367–380. [Online]. Available: http://www.usenix.org/events/osdi06/tech/madhyastha.html

[94] C. S. Leberknight, M. Chiang, H. V. Poor, and F. Wong, "A taxonomy of Internet censorship and anti-censorship," in *Fifth International Conference on Fun with Algorithms*, 2010.

[95] "How the Great Firewall of China is Blocking Tor," in *Presented as part of the 2nd USENIX Workshop on Free and Open Communications on the Internet*. Berkeley, CA: USENIX, 2012. [Online]. Available: https://www.usenix.org/conference/foci12/workshop-program/presentation/Winter

[96] Anonymous, "The Collateral Damage of Internet Censorship by DNS Injection," *SIGCOMM Comput. Commun. Rev.*, vol. 42, no. 3, pp. 21–27, Jun. 2012. [Online]. Available: http://doi.acm.org/10.1145/2317307.2317311

[97] R. Singh, H. Koo, N. Miramirkhani, F. Mirhaj, P. Gill, and L. Akoglu, "The politics of routing: Investigating the relationship between as connectivity and Internet freedom," in *6th USENIX Workshop on Free and Open Communications on the Internet (FOCI 16)*. USENIX Association, 2016.

[98] "Censorship in India by Freedom House," https://freedomhouse.org/report-types/freedom-press.

[99] J. Qiu and L. Gao, "As path inference by exploiting known as paths," 2005.

[100] "Alexa - Actionable Analytics for the Web," http://www.alexa.com/.

[101] "IP to AS Mapping, Team Cymru," http://www.team-cymru.org/IP-ASN-mapping.html.

[102] "Midar," http://www.caida.org/tools/measurement/midar/.

[103] S. Goldberg, M. Schapira, P. Hummon, and J. Rexford, "How secure are secure interdomain routing protocols," in *ACM SIGCOMM Computer Communication Review*, vol. 40, no. 4. ACM, 2010, pp. 87–98.

[104] X. Hu and Z. M. Mao, "Accurate real-time identification of IP prefix hijacking," in *Security and Privacy, 2007. SP'07. IEEE Symposium on*. IEEE, 2007, pp. 3–17.

[105] J. Qiu, L. Gao, S. Ranjan, and A. Nucci, "Detecting bogus BGP route information: Going beyond prefix hijacking," in *Security and Privacy in Communications Networks and the Workshops, 2007. SecureComm 2007. Third International Conference on*. IEEE, 2007, pp. 381–390.

[106] K. Butler, T. R. Farley, P. McDaniel, and J. Rexford, "A survey of BGP security issues and solutions," *Proceedings of the IEEE*, vol. 98, no. 1, pp. 100–122, 2010.

[107] "Pakistan hijacks Youtube." [Online]. Available: https://www.ripe.net/publications/news/industry-developments/youtube-hijacking-a-ripe-ncc-ris-case-study

[108] "Isp of oman suffers web filtering by Indian censorship," https://citizenlab.org/2012/07/routing-gone-wild/.

[109] "University of colorado, ft. collins, co, usa, "bgpmon"," http://bgpmon.netsec.colostate.edu.

[110] "Ripe ncc, amsterdam, the netherlands, "ripe ncc routing information service"," http://www.ripe.net/data-tools/stats/ris/routing-information-service.

[111] "Packet clearing house, san francisco, ca, usa," http://www.pch.net.

[112] M. Luckie, "Spurious routes in public BGP data," *ACM SIGCOMM Computer Communication Review*, vol. 44, no. 3, pp. 14–21, 2014.

[113] "Government of India Department of Telecom. Telecom Annual Report - India, 2012 - 2013, 2013." goo.gl/H7O13n.

[114] H. Acharya, S. Chakravarty, and D. Gosain, "Mending wall: On the implementation of censorship in India," in *SecureComm 2018 - 13th EAI International Conference on Security and Privacy in Communication Networks*. Springer, 2017.

[115] "List of potentially blocked websites in India — Citizen Lab," https://github.com/citizenlab/test-lists/blob/master/lists/in.csv.

[116] "Ooni: Open observatory of network interference," in *Presented as part of the 2nd USENIX Workshop on Free and Open Communications on the Internet*. Berkeley, CA: USENIX, 2012.

[117] R. Ensafi, D. Fifield, P. Winter, N. Feamster, N. Weaver, and V. Paxson, "Examining how the Great Firewall discovers hidden circumvention servers," in *Proceedings of the 2015 Internet Measurement Conference*. ACM, 2015, pp. 445–458.

[118] R. Ensafi, P. Winter, A. Mueen, and J. R. Crandall, "Analyzing the Great Firewall of China over space and time," *Proceedings on privacy enhancing technologies*, vol. 2015, no. 1, pp. 61–76, 2015.

[119] P. Pearce, B. Jones, F. Li, R. Ensafi, N. Feamster, N. Weaver, and V. Paxson, "Global measurement of DNS manipulation," in *Proceedings of the 26th USENIX Security Symposium (Security'17)*, 2017.

[120] "Internet censorship in iran." [Online]. Available: https://http://iranmediaresearch.org/en/research/pdffile/1296

[121] Z. Wang, Y. Cao, Z. Qian, C. Song, and S. V. Krishnamurthy, "Your state is not mine: a closer look at evading stateful Internet censorship," in *Proceedings of the 2017 Internet Measurement Conference*. ACM, 2017, pp. 114–127.

[122] A. Chaabane, T. Chen, M. Cunche, E. De Cristofaro, A. Friedman, and M. A. Kaafar, "Censorship in the wild: Analyzing Internet filtering in syria," in *Proceedings of the 2014 Conference on Internet Measurement Conference*. ACM, 2014, pp. 285–298.

[123] "Instances of Internet censorship in India," https://opennet.net/research/profiles/india.

[124] "Rediff," http://www.rediff.com/computer/1999/jul05dawn.htm.

[125] D. M. West, "Internet shutdowns cost countries $2.4 billion last year," *Center for Technological Innovation at Brookings, Washington, DC*, 2016.

[126] "DNA India," http://www.dnaindia.com/india/report-government-orders-blocking-of\-857-pornographic-websites-2110545.

[127] "Facebook transparency report 2017," https://transparency.facebook.com/country/India/2017-H1/.

[128] "Google transparency report 2017," https://transparencyreport.google.com/government-removals/by-country/IN.

[129] G. Aceto, A. Montieri, and A. Pescapé, "Internet censorship in italy: An analysis of 3g/4g networks," in *Communications (ICC), 2017 IEEE International Conference on*. IEEE, 2017, pp. 1–6.

[130] V. Ververis, G. Kargiotakis, A. Filasto, B. Fabian, and A. Alexandros, "Understanding Internet censorship policy: The case of greece," in *5th USENIX Workshop on Free and Open Communications on the Internet (FOCI)*, 2015.

[131] A. Dainotti, C. Squarcella, E. Aben, K. C. Claffy, M. Chiesa, M. Russo, and A. Pescapé, "Analysis of country-wide Internet outages caused by censorship," in *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference*. ACM, 2011, pp. 1–18.

[132] P. Gill, M. Crete-Nishihata, J. Dalek, S. Goldberg, A. Senft, and G. Wiseman, "Characterizing web censorship worldwide: Another look at the opennet initiative data," *ACM Transactions on the Web (TWEB)*, vol. 9, no. 1, p. 4, 2015.

[133] "Shodan-search engine for Internet-connected devices," https://www.shodan.io/.

[134] P. Pearce, R. Ensafi, F. Li, N. Feamster, and V. Paxson, "Augur: Internet-wide detection of connectivity disruptions," in *Security and Privacy (SP), 2017 IEEE Symposium on*. IEEE, 2017, pp. 427–443.

[135] Q. Yang and Y. Liu, "What's on the other side of the Great Firewall? Chinese web users' motivations for bypassing the Internet censorship," *Computers in human behavior*, vol. 37, pp. 249–257, 2014.

[136] B. Marczak, N. Weaver, J. Dalek, R. Ensafi, D. Fifield, S. McKune, A. Rey, J. Scott-Railton, R. Deibert, and V. Paxson, "China's great cannon," *Citizen Lab*, vol. 10, 2015.

[137] J. Knockel, L. Ruan, and M. Crete-Nishihata, "Measuring decentralization of Chinese keyword censorship via mobile games," in *7th {USENIX} Workshop on Free and Open Communications on the Internet ({FOCI} 17)*. USENIX Association, 2017.

[138] S. Khattak, M. Javed, P. D. Anderson, and V. Paxson, "Towards illuminating a censorship monitor's model to facilitate evasion." in *FOCI*, 2013.

[139] "Ministry of it orders isp to ban sexual abuse material," http://www.meity.gov.in/content/order-issued-meity-isps-adopt-and-implement\-iwf-resources-prevent-distribution-and.

[140] "Ooni source code," https://github.com/TheTorProject/ooni-probe/blob/master/ooni/nettests/blocking/web_connectivity.py.

[141] S. Son and V. Shmatikov, "The hitchhiker's guide to DNS cache poisoning," in *International Conference on Security and Privacy in Communication Systems*. Springer, 2010, pp. 466–483.

[142] "Bogon ip addresses," https://ipinfo.io/bogon.

[143] "Cidr report," https://www.cidr-report.org/as2.0/.

[144] L. Gao and F. Wang, "The extent of as path inflation by routing policies," in *Global Telecommunications Conference, 2002. GLOBECOM'02. IEEE*, vol. 3. IEEE, 2002, pp. 2180–2184.

[145] H. Acharya, S. Chakravarty, and D. Gosain, "Few throats to choke: On the current structure of the Internet," in *Local Computer Networks (LCN), 2017 IEEE 42nd Conference on*. IEEE, 2017, pp. 339–346.

[146] J. Jermyn and N. Weaver, "Autosonda: Discovering rules and triggers of censorship devices," in *7th USENIX Workshop on Free and Open Communications on the Internet (FOCI 17). USENIX Association, Vancouver, BC. https://www. usenix. org/conference/foci17/workshop-program/presentation/jermyn*, 2017.

[147] "Http 1.1 rfc 2616," https://tools.ietf.org/html/rfc2616.

[148] "How ooni detects http filtering?" https://ooni.torproject.org/nettest/web-connectivity/.

[149] J. Karlin, D. Ellard, A. W. Jackson, C. E. Jones, G. Lauer, D. Mankins, and W. T. Strayer, "Decoy routing: Toward unblockable Internet communication." in *FOCI*, 2011.

[150] A. Houmansadr, G. T. Nguyen, M. Caesar, and N. Borisov, "Cirripede: Circumvention infrastructure using router redirection with plausible deniability," in *Proceedings of the*

18th ACM Conference on Computer and Communications Security, ser. CCS '11. New York, NY, USA: ACM, 2011, pp. 187–200.

[151] E. Wustrow, S. Wolchok, I. Goldberg, and J. A. Halderman, "Telex: Anticensorship in the network infrastructure." in *USENIX Security Symposium*, 2011.

[152] E. Wustrow, C. M. Swanson, and J. A. Halderman, "Tapdance: End-to-middle anticensorship without flow blocking," in *23rd USENIX Security Symposium (USENIX Security 14)*, 2014, pp. 159–174.

[153] D. Ellard, C. Jones, V. Manfredi, W. T. Strayer, B. Thapa, M. V. Welie, and A. Jackson, "Rebound: Decoy routing on asymmetric routes via error messages," in *Local Computer Networks (LCN), 2015 IEEE 40th Conference on*, Oct 2015, pp. 91–99.

[154] M. Schuchard, J. Geddes, C. Thompson, and N. Hopper, "Routing around decoys," in *Proceedings of the 2012 ACM conference on Computer and communications security*. ACM, 2012, pp. 85–96.

[155] M. Nasr and A. Houmansadr, "Game of decoys: Optimal decoy routing through game theory," in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '16. New York, NY, USA: ACM, 2016, pp. 1727–1738.

[156] N. Hu and P. Steenkiste, "Exploiting Internet route sharing for large scale available bandwidth estimation," in *Proceedings of the 5th ACM SIGCOMM conference on Internet Measurement*. USENIX Association, 2005, pp. 16–16.

[157] D. Magoni and M. Hoerdt, "Internet core topology mapping and analysis," *Computer Communications*, vol. 28, no. 5, pp. 494–506, 2005.

[158] C. Orsini, E. Gregori, L. Lenzini, and D. Krioukov, "Evolution of the Internet k-dense structure," *IEEE/ACM Transactions on Networking (TON)*, vol. 22, no. 6, pp. 1769–1780, 2014.

[159] H. J. Chao and B. Liu, *High performance switches and routers*. John Wiley & Sons, 2007.

[160] S. Das, G. Parulkar, and N. McKeown, "Rethinking ip core networks," *Journal of Optical Communications and Networking*, vol. 5, no. 12, pp. 1431–1442, 2013.

[161] A. Chaabane, P. Manils, and M. A. Kaafar, "Digging into anonymous traffic: A deep analysis of the tor anonymizing network," in *Network and System Security (NSS), 2010 4th International Conference on*. IEEE, 2010, pp. 167–174.

[162] S. Dharmapurikar, P. Krishnamurthy, T. Sproull, and J. Lockwood, "Deep packet inspection using parallel bloom filters," in *High performance interconnects, 2003. proceedings. 11th symposium on*. IEEE, 2003, pp. 44–51.

[163] H. M. Moghaddam, B. Li, M. Derakhshani, and I. Goldberg, "Skypemorph: Protocol obfuscation for Tor bridges," in *Proceedings of the 19th ACM conference on Computer and Communications Security (CCS 2012)*, October 2012.

[164] Z. Weinberg, J. Wang, V. Yegneswaran, L. Briesemeister, S. Cheung, F. Wang, and D. Boneh, "StegoTorus: A camouflage proxy for the Tor anonymity system," in *Proceedings of the 19th ACM conference on Computer and Communications Security (CCS 2012)*, October 2012.

[165] A. Houmansadr, C. Brubaker, and V. Shmatikov, "The parrot is dead: Observing unobservable network communications," in *Proceedings of the 2013 IEEE Symposium on Security and Privacy*, May 2013.

[166] D. Kim, G. R. Frye, S.-S. Kwon, H. J. Chang, and A. O. Tokuta, "On combinatoric approach to circumvent Internet censorship using decoy routers," in *Military Communications Conference, MILCOM 2013-2013 IEEE*. IEEE, 2013, pp. 593–598.

[167] L. Gao, "On Inferring Autonomous System Relationships in the Internet," *IEEE/ACM Transactions on Networking (ToN)*, vol. 9, no. 6, pp. 733–745, 2001.

[168] A. Barton and M. Wright, "Denasa: Destination-naive as-awareness in anonymous communications," in *Proceedings of the 16th Privacy Enhancing Technologies Symposium (PETS 2016)*, July 2016.

[169] M. Edman and P. F. Syverson, "AS-awareness in Tor path selection," in *Proceedings of the 2009 ACM Conference on Computer and Communications Security, CCS 2009*, E. Al-Shaer, S. Jha, and A. D. Keromytis, Eds. ACM, November 2009, pp. 380–389.

[170] S. Pandey, M.-J. Choi, S.-J. Lee, and J. W. Hong, "Ip network topology discovery using snmp," in *Proceedings of the 23rd International Conference on Information Networking*, ser. ICOIN'09. Piscataway, NJ, USA: IEEE Press, 2009, pp. 33–37.

[171] "Traceroute Looking Glass," http://traceroute.org/.

[172] "IP to ASN Mapping," http://www.team-cymru.org/IP-ASN-mapping.html.

[173] "Freedom House - Freedom of Press," https://freedomhouse.org/.

[174] "Open net initiative," https://opennet.net/.

[175] "Alexa - actionable analytics for the web," http://www.alexa.com/topsites.

[176] T. D. Gautheir, "Detecting trends using spearman's rank correlation coefficient," *Environmental forensics*, vol. 2, no. 4, pp. 359–362, 2001.

[177] M. Luckie, A. Dhamdhere, B. Huffaker, D. Clark, and k. claffy, "Bdrmap: Inference of borders between ip networks," in *Proceedings of the 2016 Internet Measurement Conference*, ser. IMC '16. New York, NY, USA: Association for Computing Machinery, 2016, p. 381–396. [Online]. Available: https://doi.org/10.1145/2987443.2987467

[178] A. Marder and J. M. Smith, "Map-it: Multipass accurate passive inferences from traceroute," in *Proceedings of the 2016 Internet Measurement Conference*, 2016, pp. 397–411.

[179] A. Marder, M. Luckie, A. Dhamdhere, B. Huffaker, k. claffy, and J. M. Smith, "Pushing the boundaries with bdrmapit: Mapping router ownership at internet scale," in *Proceedings of the Internet Measurement Conference 2018*, ser. IMC '18. New York, NY, USA: Association for Computing Machinery, 2018, p. 56–69. [Online]. Available: https://doi.org/10.1145/3278532.3278538

[180] M. Luckie, B. Huffaker, A. Dhamdhere, V. Giotsas *et al.*, "As relationships, customer cones, and validation," in *Proceedings of the 2013 conference on Internet measurement conference*. ACM, 2013, pp. 243–256.

[181] D. Gosain, A. Agarwal, S. Chakravarty, and H. Acharya, "The devil's in the details: Placing decoy routers in the Internet," in *Proceedings of the 33rd Annual Computer Security Applications Conference*. ACM, 2017, pp. 577–589.

[182] A. Barton and M. Wright, "Denasa: Destination-naive as-awareness in anonymous communications," *Proceedings on Privacy Enhancing Technologies*, vol. 2016, no. 4, pp. 356–372, 2016.

[183] "Google global cache," https://peering.google.com/#/.

[184] B. Yeganeh, R. Rejaie, and W. Willinger, "A view from the edge: A stub-as perspective of traffic localization and its implications," in *2017 Network Traffic Measurement and Analysis Conference (TMA)*. IEEE, 2017, pp. 1–9.

[185] R. Deibert and R. Rohozinski, "Liberation vs. control: The future of cyberspace," *Journal of Democracy*, vol. 21, no. 4, pp. 43–57, 2010.

[186] J. Holowczak and A. Houmansadr, "Cachebrowser: Bypassing Chinese censorship without proxies using cached content," in *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2015, pp. 70–83.

[187] D. Fifield, C. Lan, R. Hynes, P. Wegmann, and V. Paxson, "Blocking-resistant communication through domain fronting," *Proceedings on Privacy Enhancing Technologies*, vol. 2015, no. 2, pp. 46–64, 2015.

[188] R. McPherson, A. Houmansadr, and V. Shmatikov, "Covertcast: Using live streaming to evade Internet censorship," *Proceedings on Privacy Enhancing Technologies*, vol. 2016, no. 3, pp. 212–225, 2016.

[189] B. Gueye, A. Ziviani, M. Crovella, and S. Fdida, "Constraint-based geolocation of Internet hosts," *IEEE/ACM Transactions on Networking (TON)*, vol. 14, no. 6, pp. 1219–1232, 2006.

[190] Z. Weinberg, S. Cho, N. Christin, V. Sekar, and P. Gill, "How to catch when proxies lie: Verifying the physical locations of network proxies with active geolocation," in *Proceedings of the Internet Measurement Conference 2018*. ACM, 2018, pp. 203–217.

[191] P. Syverson, R. Dingledine, and N. Mathewson, "Tor: The second generation onion router," in *Usenix Security*, 2004.

[192] D. Cicalese, D. Giordano, A. Finamore, M. Mellia, M. Munafò, D. Rossi, and D. Joumblatt, "A first look at anycast cdn traffic," *arXiv preprint arXiv:1505.00946*, 2015.

[193] F. Wohlfart, N. Chatzis, C. Dabanoglu, G. Carle, and W. Willinger, "Leveraging interconnections for performance: The serving infrastructure of a large cdn," in *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication*, ser. SIGCOMM '18. New York, NY, USA: ACM, 2018, pp. 206–220. [Online]. Available: http://doi.acm.org/10.1145/3230543.3230576

[194] X. Fan, E. Katz-Bassett, and J. Heidemann, "Assessing affinity between users and cdn sites," in *International Workshop on Traffic Monitoring and Analysis*. Springer, 2015, pp. 95–110.

[195] M. Calder, A. Flavel, E. Katz-Bassett, R. Mahajan, and J. Padhye, "Analyzing the performance of an anycast cdn," in *Proceedings of the 2015 Internet Measurement Conference*. ACM, 2015, pp. 531–537.

[196] E. Bos, "Analyzing the performance of cloudflare anycast cdn, a case study," in *27th Twente Student Conference on IT*, 2017.

[197] B. Ager, W. Mühlbauer, G. Smaragdakis, and S. Uhlig, "Web Content Cartography," in *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference*. ACM, 2011, pp. 585–600.

[198] M. Gharaibeh, A. Shah, B. Huffaker, H. Zhang, R. Ensafi, and C. Papadopoulos, "A look at router geolocation in public and commercial databases," in *Proceedings of the 2017 Internet Measurement Conference*. ACM, 2017, pp. 463–469.

[199] M. Calder, X. Fan, Z. Hu, E. Katz-Bassett, J. Heidemann, and R. Govindan, "Mapping the expansion of google's serving infrastructure," in *Proceedings of the 2013 conference on Internet measurement conference*. ACM, 2013, pp. 313–326.

[200] T. Böttger, F. Cuadrado, G. Tyson, I. Castro, and S. Uhlig, "Open connect everywhere: A glimpse at the Internet ecosystem through the lens of the netflix cdn," *ACM SIGCOMM Computer Communication Review*, vol. 48, no. 1, pp. 28–34, 2018.

[201] J. Xue, D. Choffnes, and J. Wang, "Cdns meet cn an empirical study of cdn deployments in China," *IEEE Access*, vol. 5, pp. 5292–5305, 2017.

[202] W. Scott, T. Anderson, T. Kohno, and A. Krishnamurthy, "Satellite: Joint analysis of cdns and network-level interference," in *2016 {USENIX} Annual Technical Conference* 16, 2016, pp. 195–208.

[203] R. Ensafi, D. Fifield, P. Winter, N. Feamster, N. Weaver, and V. Paxson, "Examining how the Great Firewall discovers hidden circumvention servers," in *Internet Measurement Conference*. ACM, 2015.

[204] T. K. Yadav, A. Sinha, D. Gosain, P. K. Sharma, and S. Chakravarty, "Where The Light Gets In: Analyzing Web Censorship Mechanisms in India," in *Proceedings of the Internet Measurement Conference 2018*. ACM, 2018, pp. 252–264.

[205] S. Aryan, H. Aryan, and J. A. Halderman, "Internet censorship in Iran: A first look," in *Presented as part of the 3rd $USENIX$ Workshop on Free and Open Communications on the Internet*, 2013.

[206] V. Ververis, G. Kargiotakis, A. Filasto, B. Fabian, and A. Alexandros, "Understanding Internet censorship policy: The case of greece," in *5th $USENIX$ Workshop on Free and Open Communications on the Internet ($FOCI$)*, 2015.

[207] M. C. Tschantz, S. Afroz, V. Paxson *et al.*, "Sok: Towards grounding censorship circumvention in empiricism," in *2016 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2016, pp. 914–933.

[208] S. Khattak, T. Elahi, L. Simon, C. M. Swanson, S. J. Murdoch, and I. Goldberg, "Sok: Making sense of censorship resistance systems," *Proceedings on Privacy Enhancing Technologies*, vol. 2016, no. 4, pp. 37–61, 2016.

[209] H. Zolfaghari and A. Houmansadr, "Practical censorship evasion leveraging content delivery networks," in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2016, pp. 1715–1726.

[210] E. Pujol, P. Richter, B. Chandrasekaran, G. Smaragdakis, A. Feldmann, B. M. Maggs, and K.-C. Ng, "Back-office web traffic on the Internet," in *Proceedings of the 2014 Conference on Internet Measurement Conference*. ACM, 2014, pp. 257–270.

[211] "Maxmind geoip2 database," https://www.maxmind.com/en/geoip2-services-and-databases.

[212] B. Chandrasekaran, M. Bai, M. Schoenfield, A. Berger, N. Caruso, G. Economou, S. Gilliss, B. Maggs, K. Moses, D. Duff *et al.*, "Alidade: Ip geolocation without active probing," *Department of Computer Science, Duke University, Tech. Rep. CS-TR-2015.001*, 2015.

[213] I. Poese, S. Uhlig, M. A. Kaafar, B. Donnet, and B. Gueye, "Ip geolocation databases: Unreliable?" *ACM SIGCOMM Computer Communication Review*, vol. 41, no. 2, pp. 53–56, 2011.

[214] V. N. Padmanabhan and L. Subramanian, "Determining the geographic location of Internet hosts," in *SIGMETRICS/Performance*, 2001, pp. 324–325.

[215] E. Katz-Bassett, J. P. John, A. Krishnamurthy, D. Wetherall, T. Anderson, and Y. Chawathe, "Towards ip geolocation using delay and topology measurements," in *Proceedings of the 6th ACM SIGCOMM conference on Internet measurement*. ACM, 2006, pp. 71–84.

[216] B. Wong, I. Stoyanov, and E. G. Sirer, "Octant: A comprehensive framework for the geolocalization of Internet hosts." in *NSDI*, vol. 7, 2007, pp. 23–23.

[217] Y. Wang, D. Burgener, M. Flores, A. Kuzmanovic, and C. Huang, "Towards street-level client-independent ip geolocation." in *NSDI*, vol. 11, 2011, pp. 27–27.

[218] S. Laki, P. Mátray, P. Hága, T. Sebők, I. Csabai, and G. Vattay, "Spotter: A model based active geolocation service," in *2011 Proceedings IEEE INFOCOM*. IEEE, 2011, pp. 3173–3181.

[219] R. Percacci and A. Vespignani, "Scale-free behavior of the Internet global performance," *The European Physical Journal B-Condensed Matter and Complex Systems*, vol. 32, no. 4, pp. 411–414, 2003.

[220] A. Dhamdhere and C. Dovrolis, "The Internet is flat: modeling the transition from a transit hierarchy to a peering mesh," in *Proceedings of the 6th International COnference*. ACM, 2010, p. 21.

[221] I. Castro, J. C. Cardona, S. Gorinsky, and P. Francois, "Remote peering: More peering without Internet flattening," in *Proceedings of the 10th ACM International on Conference on emerging Networking Experiments and Technologies*. ACM, 2014, pp. 185–198.

[222] "Citizen Lab list of blocked websites," https://github.com/citizenlab/test-lists.

[223] J. A. Suykens and J. Vandewalle, "Least squares support vector machine classifiers," *Neural processing letters*, vol. 9, no. 3, pp. 293–300, 1999.

[224] R. Kohavi *et al.*, "A study of cross-validation and bootstrap for accuracy estimation and model selection," in *Ijcai*, vol. 14, no. 2. Montreal, Canada, 1995, pp. 1137–1145.

[225] B. Huffaker, M. Fomenkov *et al.*, "Geocompare: a comparison of public and commercial geolocation databases-technical report," Cooperative Association for Internet Data Analysis (CAIDA), Tech. Rep., 2011.

[226] "Ripe atlas credit system," https://atlas.ripe.net/docs/credits/.

[227] "Cisco Umbrella Popular Websites," https://umbrella.cisco.com/blog/2016/12/14/cisco-umbrella-1-million/.

[228] "mmjestic Million Popular Websites," https://majestic.com/reports/majestic-million.

[229] "Pluggable transport for Tor," https://2019.www.torproject.org/docs/pluggable-transports.html.en.

Appendices

.1 Path frequency vs customer-cone size

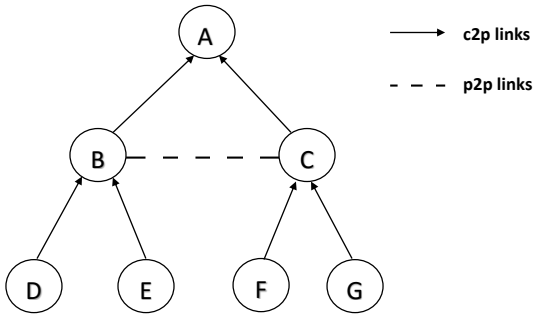

Figure 1: Schematic AS graph with multiple valid valley-free paths: $D-B-E$, $D-B-C-F$, $D-B-C-G$, $D-B-A-C-F$, $D-B-A-C-G$, $E-B-A-C-F$ and $E-B-A-C-G$. Note how some do not traverse A, the AS with the highest customer-cone size.

We provide some detail for our claim in Section 6.4, that customer-cone size is not a reliable metric to identify the ASes that transport a large fraction of traffic. We explain our reasoning with the example of the AS graph in Figure 1.

The figure represents a hypothetical AS graph where node A represents an AS with the highest customer-cone size of 6, the total number of ASes that A can reach via its customers and their customers (D, B, E, F, C, G). ASes B and C have customer cones of size 2 (for each of the individual nodes).

There are several valid valley free paths in this hypothetical AS graph: $D-B-E$, $D-B-C-F$, $D-B-C-G$, $D-B-A-C-F$, $D-B-A-C-G$, $E-B-A-C-F$ and $E-B-A-C-G$. However, as evident from the example, not all of them pass through the *root AS*, *i.e.* the node with the highest customer-cone size. This is also evident from Table 1: for several large ASes, a considerable fraction of paths do not traverse the core ASes themselves, but do traverse their immediate (1-hop) customers [1].

In fact, customer cone size is not even well correlated with path frequency. The Spearman's Rank Correlation Coefficient is only ≈ 0.2 (see Figure 2).

[1] Interestingly, smaller customer-cones have fewer such paths, i.e. paths not traversing the root AS. Perhaps smaller cones have fewer neighbor ASes to route through? We may study this in future work

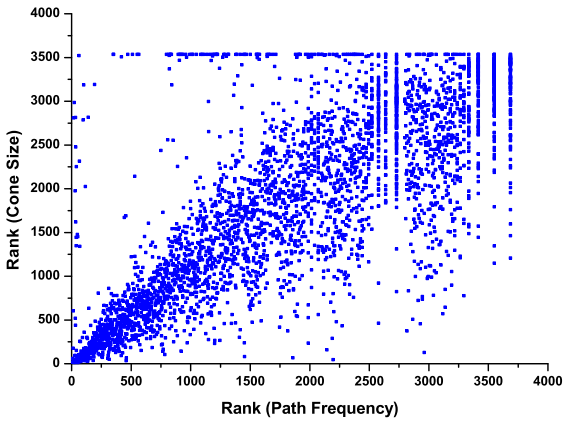

Figure 2: AS Rank variation: path frequency vs cone size for transit ASes.

.2 Additional Graphs

Traffic to specific destinations:

For completeness, we present the results of our experiment for a few of the most important single websites in isolation, in Figure 3.

Clearly, while single websites are far more variable, the general trend is similar to Figure 5.3. About 15 ASes (out of the 50 heavy-hitter ASes identified) cover over 80% of the AS-paths to

ASN	% of path not reaching the AS	% of path reaching the AS
3356	34.16	33.17
174	29.05	13.13
2914	28.16	12.90
1299	36.50	8.05
3257	21.00	5.23
6939	7.46	4.40
6461	5.13	4.03
6453	26.00	3.76
7018	7.40	3.70
10310	0.07	3.52

Table 1: Prefix-to-AS paths in cone of core ASes: %age traversing core ASes themselves, vs. %age traversing their immediate (1-hop) customers.

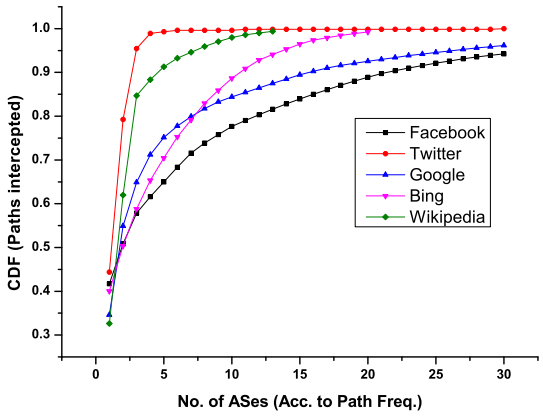

Figure 3: CDF of ASes according to fraction of paths to popular websites that they intercept

these destinations.

Only 5 ASes collectively transport all the paths to the prefix corresponding to `twitter.com`, while about 18 ASes intercept all paths carrying traffic from `bing.com`. Finally, about 30 ASes cumulatively transport about 98% of the paths corresponding to `google.com` and `facebook.com`.

For comparison, here is the cumulative frequency of the paths intercepted by ASes, this time for paths to the Alexa top-200 sites.

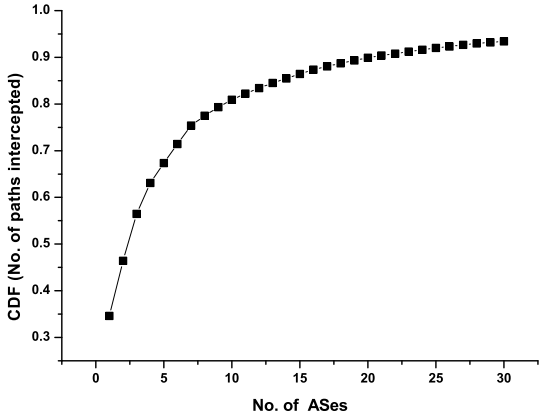

Figure 4: CDF of ASes by fraction of paths that they intercept (for Alexa top-200 sites)

We also provide the graph of the cumulative path coverage inside an AS, by its heavy hitter routers. As is quite clear, the graph varies a good deal; some ASes are almost completely covered

by a few routers, while others (AS 3257) need very many heavy-hitters, and can more easily be covered by choosing their edge routers (Figure 5).

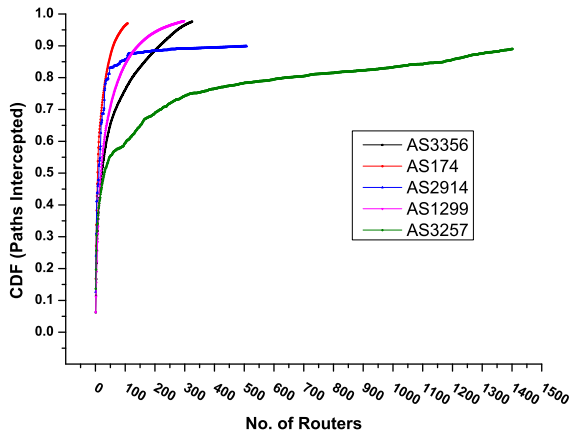

Figure 5: `Traceroute` paths in the top five (out of 30) key ASes. The number of routers needed to cover 90% of the paths varies between 288 (AS174) and 1483 (AS3257)

.3 RTT Profiles for RIPE Probes in Various Countries

Figures 6, 7, 8, 9 represent the SVM line (threshold) for most consistent reference nodes in India, Saudi Arabia, Brazil and Unites States.

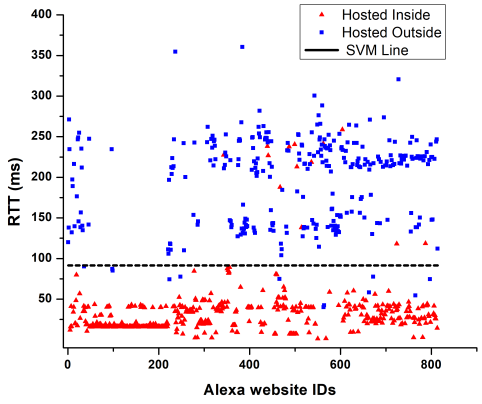

Figure 6: RTT scatter plot for a probe in India.

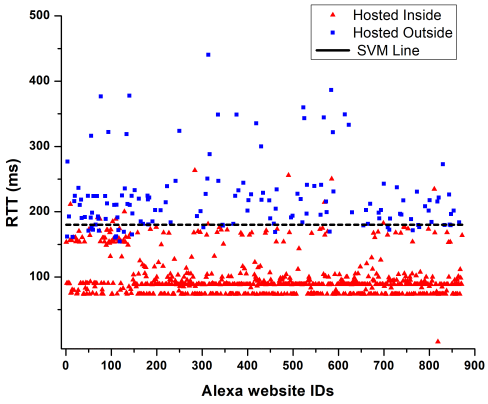

Figure 7: RTT scatter plot for a probe in Saudi Arabia.

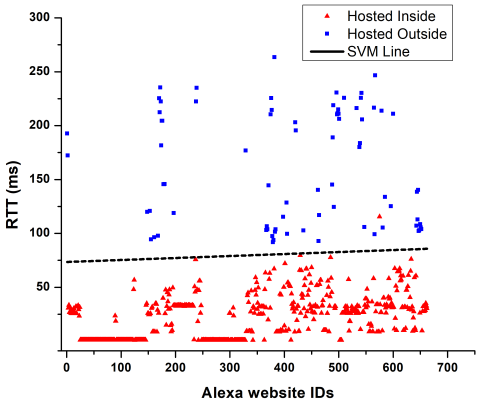

Figure 8: RTT scatter plot for a probe in Brazil.

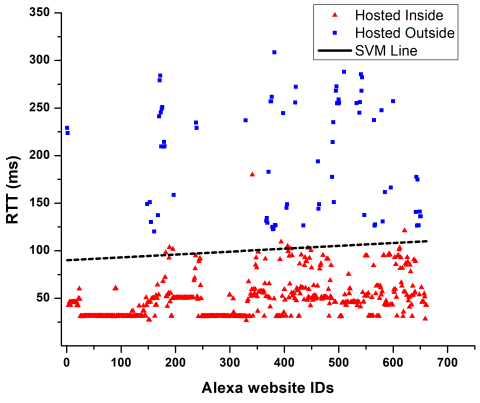

Figure 9: RTT scatter plot for a probe in United States.