Cairn:Identifying Network Locations for Large Scale Censorship by Resource-Constrained Adversaries

Student Name: Anshika Agarwal

IIIT-D-MTech-CS-GEN-MT14048
June, 2016

Indraprastha Institute of Information Technology
New Delhi

Thesis Committee
Dr. Sambuddho Chakravarty (Advisor)
Dr. Pushpendra Singh (Internal Examiner)
Dr.Anupam Joshi(External Examiner)

Submitted in partial fulfillment of the requirements
for the Degree of M.Tech. in Computer Science,
in General Category

Keywords: Internet Censorship, autonomous system,traceroute,pefix hijacking, collateral damage,network topology and dns resolvers

Certificate

This is to certify that the thesis titled **"Cairn:Identifying Network Locations for Large Scale Censorship by Resource- Constrained Adversaries"** submitted by **Anshika Agarwal** for the partial fulfillment of the requirements for the degree of *Master of Technology* in *Computer Science & Engineering* is a record of the bonafide work carried out by her / him under my / our guidance and supervision in the Security and Privacy group at Indraprastha Institute of Information Technology, Delhi. This work has not been submitted anywhere else for the reward of any other degree.

Dr.Sambuddho Chakravarty
Indraprastha Institute of Information Technology, New Delhi

Abstract

Censorship Of the Internet by government is a hotly contested topic.Some nations lean more towards free speech;others are much more conservative. How feasible is it for a government to censor the Internet? What mechanisms can it use? Where all shall it install the censorship infrastructure? What collateral damage can be seen in other countries? In this paper,we attempt to look at these questions in general, and present a case study of India-a country which currently performs limited censorship,but which will likely change its access policies in the near future.

Acknowledgments

I am extremely grateful and remain indebted to my guide Dr.Sambuddho Chakravarty for being a source of inspiration and for his constant support in the Design,Implementation and Evaluation of the thesis.The quality of the work would not have been as high without his well-appreciated advice.It is an honor and privilege to have him as my mentor.I thank him from the bottom of my heart.I thank Dr.H.B Acharya for enriching this thesis with his valuable suggestions and feedback. Special thanks to Devashish Gosain for critically reviewing and commenting on the work and helping to make it what it is. I would like to thank all the members of the CERC(Cybersecurity Education and Research Centre) for their valuable suggestions.

My heartfelt thanks to Rahul Singh for his consistent guidance and valuable time which he has given from his busy schedule. Special thanks to my friends Ashish Bandil and Ajit Pratap Singh for helping me throughout this thesis.I would like to thank my parents for always supporting me and providing me with an environment where I could work dedicatedly.

Last but not the least,I would like to thank IIITD for making this happen.The infrastructure,services and environment provided to us as students are truly remarkable.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Censoring the Internet is a hard problem: it is decentralized,routes around disruptions,and is designed to be resilient to disruptions.While a sufficiently resourceful adversary can filter(or monitor) almost all traffic [16],this is not true for resource constrained nations,who cannot afford to place dedicated infrastructure in every network.But exactly how hard is it for a given government to become censorious? can it,for example,make do with network monitors in those few key Autonomous Systems(ASes) and network elements that forward most of the nation's traffic? How can we measure the potential censoring power of a regime? (The US censors much less than,for example,Iran. But if it did decide to censor content, the effect would be much greater).How much collateral damage will a nation cause, if it becomes censorious? In this thesis, we make a start toward answering such questions.

The case we study is India, a democratic nation that favors free speech, but undecided on questions of censorship [4].(For example, in August 2015, Indian ISPs blocked 857 pornographic and torrent sites under State orders [9]; public outcry forced them to backtrack.) There are several known engineering challenges in implementing a censorship scheme; for example, how to identify the content to block, *etc.* We take a new approach, and focus on the network itself - which ASes and network elements would be most effective for installing censorship infrastructure, how effective it would be, and what collateral damage would result. Formally, we aim to answer the following questions.

- *Is it feasible to filter or monitor Internet traffic? If so, how, and where?* While India is among top Internet consumers, with over 300 million users, the Internet in India consists of only about 900 active Autonomous Systems (ASes), of which less than 100 are ISP networks (the others are stubs, relying on the ISPs for Internet connectivity). Is the Internet in India sufficiently localized that it can be effectively policed, or are there too many "throats to choke"? What mechanisms might be effective
  - Is there a small number of key ASes (networks) and routers, such that the government can deploy network monitoring or filtering infrastructure at only these ASes and intercept most of the traffic to censored destinations?
  - Which censorious ISPs deploy DNS injectors [14] so that they can filter most DNS requests?
  - What fraction of Indian ASes would be affected if censorious ISPs choose to hijack IP prefixes [7]?
- *Will traffic filtering cause collateral damage?:* Some nations get Internet connectivity through others; as a consequence, censorship by an "upstream" nation can lead to collateral

damage, where requests for content originating from the "downstream" nation are forcibly filtered. If India becomes a censorious regime, to what extent will this affect traffic of non-Indian origin?

1.1 Problem Description

The structure of the current commercial Internet with its provider/customer relation between Autonomous System makes it vulnerable to adversary who controls few key ASes only.Here, we explore the power of an adversary that wishes to censor traffic in any large country. Given that most governments do not have the resources for an extensive monitoring solution like PRISM [22], the question arises how powerful a resource-constrained governmental adversary is. Specifically, we are interested in the following questions:

- Is it possible for the Government to monitor/censor a large fraction of Internet traffic by controlling only a small number of network locations? And by what means?
- What fraction of traffic could be filtered, and who would be most affected?
- Would such censorship affect users outside the country as well?

In this thesis, we explore these questions with a case study of India.

Threat Model: Our adversary is a censorious but resource-budgeted government. The adversary aims to monitor and filter Internet traffic, and for this purpose may perform IP filtering, DNS injection, and IP prefix hijacking attacks.

1.2 Motivation

According to ONI, India is engaged in selective filtering and Freedom House also declares India to be a partly free regime. It is a democratic regime and time and again central/state government has ordered to put a ban on many websites (including pornographic content). Since India is a resource constrained nation and there is a lacuna of sophisticated techniques like DPI enabled routers, it becomes imperative to find some Key places where if censorship techniques employed, can cause massive affect.

The problem in this thesis is to determine where in the resource-constrained country, government should deploy network monitoring or filtering infrastructure in order to have maximum impact from less number of resources.This impact is measured by the fraction of traffic paths passing through an AS (and the routers within it).We aim to find some key ASes (networks and routers) which intercepts maximum traffic.

A sufficiently resourceful adversary can filter (or monitor) almost all traffic.But it is not true for a resource constrained nation. It does not have enough resources to deploy firewall everywhere

in order to carry out censorship. For such nations , we aim to find some key places in the network map which will intercept maximum traffic.

In order to find the key routers,i.e. a small number of routers that carry most of the Internet traffic , we proceed in next two phases.First,we identify the key ASes in the Internet; next, we identify the key routers in these ASes.

In order to find key ASes , map of the Internet is constructed, and select the ASes that occur most frequently in our paths (taken from BGP routing tables). In the next phase,top five ASes (which intercepts most of the traffic) are mapped, to identify their key routers; this allows us to estimate the number of firewalls/network infrastructure we need to be able to intercept a large fraction of Internet traffic.

1.3 Background

The interaction of the Internet with government policy (especially censorship and privacy issues) is an extensively studied subject. Government interference with the Internet comes under two main headings, as follows.

1. Network censorship, where specific traffic is filtered and blocked.
2. Network surveillance - silently monitoring traffic. (This is harder to detect.)

In this thesis, we study the structure of the network in India (i.e., the connectivity of autonomous systems, and how they forward traffic). The aim is to identify key points through which almost all of the traffic, DNS requests, *etc.* must pass. *These key points would be the most cost-effective locations to perform censorship or surveillance.*

This thesis relies heavily on mapping the structure of the Internet, an area of research called *network tomography* [18]. 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). There are over $40,000$ ASes in the Internet, including ISPs (e.g. AT&T) and customer networks. Consequently, Internet mapping proceeds at two levels:

1. *AS-level mapping.* Gao *et al.* [23] show how to infer the paths from a given IP address prefix to every AS on the Internet. The algorithm uses publicly-available BGP routes, obtained from various Internet Exchange Points across the globe [12]), to estimate the relationships and connections between ASes, and builds a directed graph of the Internet where each node is an AS.
2. *Router-level mapping.* An AS is not a black box, but contains hosts and routers. Mahajan *et al.* [25] show how the internal structure of an AS can be mapped, by a combination of `traceroute` probes, IP alias resolution [1], and reverse DNS lookups.

[1]Different interfaces of the same router, with different IP addresses, are called IP aliases

1.3.1 Network censorship

Internet Censorship is defined as 'refusing users access to certain web pages without the cooperation of the content provider, the hosting provider and the owner of the client machine being used to access these pages. Internet censorship is encouraged by many countries as they do not want their citizens to be influenced by western ideology and western culture. They do not want their citizens to access content related to pornography, violence and other illegal content. Open Net Initiative mentioned four nodes in the network where Internet Filtering can occur [3]: **Internet Backbone** - All the blocking technologies may be carried out at the backbone level, which affects Internet access throughout the entire country. This is often carried out at the International gateway. **Internet Service Providers** - Censorship which is done at government level is mostly carried out by Internet Service Providers by using any of the technical methods used for filtering. **Institutions** - Institutional level filtering occurs in various organizations, colleges, schools and government companies. This type of filtering is mainly carried out to meet the requirements of the institution. Various colleges censor Facebook and other social media sites so that students may focus on studies. Sometimes, government also impose censorship on some institutions. **Individual Computers** - Home level filtering can be achieved by installation of filtering software that restricts an individual's computer ability to access certain sites.

Technical Approaches to censorship

Internet is based on Client Server model and so censorship can be classified as Client based censorship (if filter is placed on the client side requesting the data), Server based censorship (if filter is placed on the server returning the data) and Network based censorship (if filter is placed on the network between client and server).

- Client-based censorship: In this approach access to online content is restricted by running censorship promoting applications on the same network to which the client is connected. For instance while typing using keyboard, if match occurs with blacklisted words, access is denied.
- Server based censorship In this approach Censorship occurs at server side. Server stealthily and selectively monitors, intercepts, redirect, modify and restricts access to some online content. This form of censorship is hard to discover as it is silently active on one end of the communication away from user.
- Network based censorship In this case censorship occurs between client and server somewhere in the network. This type of censorship is a trade-off between the pros and cons of two aforementioned techniques. With Client based adversary has wide coverage but more chances of possible detection as application is running on client machine. With Server based it is completely vice-versa. Network based stands between them.

There are majorly three technical methods to block access to Internet sites as **IP Filtering, DNS Tampering and URL blocking**. These techniques are used to block access to specific web pages, domains or IP address. Keyword Blocking is the most advanced technique which is growing nowadays and carried out by various countries. It blocks access to websites based on the keyword found in URL which matches blacklisted list of words.

This thesis aims to find optimal locations to employ these techniques by a resource constrained adversary like India.

1.3.2 Network tomography

Network tomography [18], the mapping of the structure of the Internet, involves several kinds of map. The Internet consists of routers and hosts, organized into networks called Autonomous Systems (ASes) that decide how to route traffic among themselves. Besides physical connections, there must be an acceptable business relationship before an AS will route traffic over a link. There is also the question of the geographic locations of networks and hosts.

For our research, we require maps representing the connections between ASes and also IP-level connectivity within these ASes.

AS-level mapping.: The approach taken in earlier work, such as the CAIDA Ark project [2], involves mapping 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.

However, the CAIDA approach simply finds all available paths, and makes no distinction between heavily-used paths and ones that only exist in theory. *The actual path between ASes may be decided by factors other than path length*[3]. We use a different approach, following Gao and Qiu [24]. This algorithm builds paths from a given IP prefix to every AS of the Internet, using publicly-available BGP routes (which we obtain from Internet Exchange Points across the globe [10]). It also allows us to infer inter-AS relationships.

Router-level mapping. An large AS, such as an ISP, generally has several thousand routers. Just as we mapped the inter-AS graph using BGP information, it is possible to map their internal structure using their SNMP Management Information Bases (MIBs) [**?**]. However, we have no access to this data. Instead, we map the routers and connections in ASes of interest using the Rocketfuel algorithm [26].

In brief, Rocketfuel runs `traceroute` probes from looking glass servers [11] to prefixes inside a chosen AS. We also resolve IP aliases [2] using MIDAR [6], and perform reverse DNS lookups. Our mapping is similar to that of CAIDA, but has much fresher data.

[2]Different interfaces of the same router, with different IP addresses, are called IP aliases

Chapter 2

Literature Survey

Censorship has been a topic of research for many social and political scientists encompassing various analog and mass media. In recent times, it has seen a flurry of activities, from computer security and network scientists. Oppressive regimes adopts various methods to perform network wide censorship. For example BGP tempering resulted in Arab spring (a mass unrest in Egypt and Libya in 2011). In these cases the countries were made unreachable from ASes outside of the country by withdrawing their presence from the BGP network view. When these techniques escape the control of the censor, dramatic side effects occur, and in 2009 accidentally youtube was made unreachable to outside world when Pakistan advertised the fake prefix of youtube.

Not only BGP tempering, but DNS has also been used to achieve censorship. For instance China uses DNS injection to block not only Chinese traffic but also traffic transiting through Chinese ASes contributing to collateral damage [19]. China has placed many DNS injectors on the paths to DNS revolvers, when DNS request to resolver passes the injector, it is suspected for possible blocked content. If packet contains request to blocked sites or has filtered keywords,it is blocked and user receives either `NXDOMAIN` or a static page.

This thesis contributes to the study of country-level network censorship. Much work in this area focuses on China; for example, Winter *et al.* [27] examine how the Chinese authorities use DPI-capable routers to detect Tor Bridges. A major step forward was made by Verkamp *et al.* [13], who deployed clients in 11 countries to identify their network censorship activities – IP and URL filtering, keyword filtering DNS censorship *etc.* He found that censorship mechanism vary across the countries. Malaysia, Russia and Turkey censor at the DNS while other countries do so at routers or other network hardware. South Korea censors both at the DNS and at the routers. Also, Saudi Arabia and China censor on destination IP addresses, while other countries primarily censor on hostnames, URLs, or keywords. Further, execution of censorship ranges between DNS redirects (Malaysia, Russia, South Korea and Turkey), connection timeouts (Bangladesh and India), TCP resets (China), and HTTP responses with various 200-, 300-, and 400-level status codes (Bahrain, Iran, Saudi Arabia, South Korea and Thailand). Later authors - Nabi [21] in Pakistan, and Halderman *et al.* [15] in Iran - demonstrate different methods of censorship employed by their respective regimes, as well as different forms of content blocked. In Pakistan,

Nabi found that sites were primarily blocked with either DNS hijacking or HTTP filtering. While in Iran more diverse form of censorship is observed using HTTP Hostbased blocking, keyword filtering, DNS hijacking, and protocol-based throttling. However, such a study of censorship in repressive regimes is necessarily limited, as it requires Internet access from inside the country. (Nabi was able to run probes from only five locations, and Halderman from only one.) Present state of art, describes the existing methods which are already used in use by various countries to perform internet censorship. We take a different direction with this thesis and study censorship as an engineering problem. If tomorrow Indian Government wants a How feasible is it for the resource constrained State to carry out censorship by standard means (DNS injection attacks, IP filtering and IP prefix hijacking)? How many, and which, autonomous systems would it need to control? We explain our approach in detail in the next Section.

Chapter 3

Proposed Methodology

3.1 Identifying Potential network locations for IP filtering

In order to estimate the locations for installing IP filtering infrastructure, we mapped the complete Internet topology consisting of ASes as nodes, then focused on Indian ASes and their connections. Gao *et. al.*'s algorithm [23] was used to find the AS paths connecting the home AS of some chosen IP prefixes (corresponding to censored sites) to every other AS in the world. (The algorithm builds a graph of ASes and their links from known AS paths in BGP routing tables, obtained from a number of vantage points [12], and known business relationships between ASes [17].)

Unlike other nations, which have an unambiguous list of blocked sites [21], India has no clear censorship policy. We chose sites reported as blocked in India, from the crowd-sourced censorship-reporting site, Herdict [3]. These included social networking sites, political sites, sites related to unfriendly nations, and p2p file-sharing sites. We also chose several adult sites from Alexa [1], popularly accessed in India. Our set of censored websites corresponded to 211 unique IP prefixes.

We then calculated the paths between all Indian ASes and these prefixes, a total of 186679 paths. The ASes appearing in these paths were sorted by frequency of occurrence; we selected the few most popular.

Intra-AS topology generation and network "choke" routers: The router level topology was computed for two major ASes (ISPs) that appear in a very large fraction of paths. Using `planetlab` nodes, we ran `traceroute` probes to one representative IP in each prefix advertised by the ASes. (Our approach follows Mahajan *et al.* [25].) We then resolved aliases with `Midar` [6], selected the routers common to many `traceroute` paths, and attempted to identify any "key" routers, a small set of which could give the adversary coverage of all paths.

3.2 Identifying Potential Sites for DNS Injection

Another common approach to censorship is to prevent the DNS service from resolving requests - the censor intercepts DNS connections and responds back with bogus IPs or NXDOMAIN responses. This is referred to as *DNS Injection*.

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` [20]. The probe requests evoke three kinds of responses: *open*, *filtered*, and *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 [1].)

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.

3.3 Impact of IP Prefix Hijack Based Censorship

IP Prefix Hijacking attack involves malicious BGP routers advertising fake AS paths, in an attempt to poison routes to an IP prefix (*e.g.* [7]). Such advertisements make the router appear to be an attractive choice for routing traffic to the prefix [16] (see Figure 3.1).

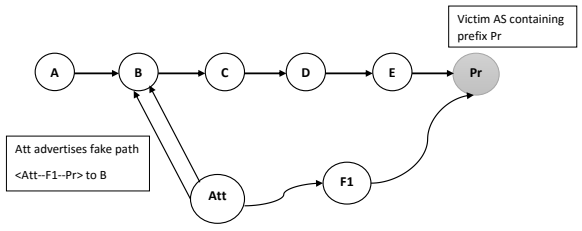

Figure 3.1: IP Prefix Hijacking: Valid path: $A-B-C-D-E-Pr$. A is the origin AS and Pr the AS with the destination prefix. Attacker Att advertises a shorter path $Att-F_1-Pr$, to AS B. If B chooses this path and directs its traffic to Att, the attacker can censor the traffic.

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. Receiving ASes accept these advertisements based on the following heuristic [16]:

1. If there exists a customer path towards the target IP and iff the advertisement presents a

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

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.

3.3.1 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*[2]. Inspecting the prefix-to-AS paths, we identified ASes with which the attacker AS had a business relationship, and applied the above heuristics to determine the number of ASes potentially affected by fake advertisements.

3.4 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. In the past, such inadvertent filtering has been reported by Sparks *et. al.* [14]. As one of our research objectives, we try to identify ASes outside India that may be affected by Indian censorship. To that end we identified 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.

3.5 Data Collection

A.Network mapping process

Our network mapping process consists of two phases. First, we build an AS-level Internet map, using the paths connecting popular WWW destinations (which are, in fact, also our potential censored sites in case of IP Filtering and DNS resolvers in case of DNS Injection) and the various ASes of the Internet. Using the resultant paths, we identify ASes of high path frequency, i.e. those that appear in a large number of paths, as key ASes (for hosting firewalls or dns injectors).

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. These are the routers that we intend to replace with firewalls or dns injectors.

[2]Number of ASes that are adjacent to the said AS

Generating AS level maps

For the first phase of network mapping, we use the approach presented by Gao *et al.* [24]. Their approach uses existing AS paths appearing in BGP tables collected from a number of Internet Exchange Points(IXes) [10], and infers paths that do not explicitly appear. Existing BGP paths are augmented by appending other ASes that frequently appear adjacent to path ASes, and which do not invalidate the path's *Valley-Free* property. The aim is to build paths connecting a given IP prefix to all ASes in the Internet.

For our analysis, we use a snapshot of BGP routes corresponding to March 1, 2016, derived by merging the Routing Information Bases (RIBs) obtained from 15 Internet Exchange Points. Taking the IP address prefixes corresponding to the top 100 most-popular sites, we generate paths connecting all ASes to these prefixes.

Finally, we rank the ASes according to the number of times they appear in different paths. (We provide details of these ASes, and the fraction of paths that they intercept, in Section **??**.

Creating router level maps

After identifying key ASes in the Internet, as above, we are still left with the problem of where in the AS to put a decoy router. A practical AS involves a complex topology of routers and hosts. We map the topology of routers inside some key ASes, so as to identify the actual routers that potentially transport large fraction of the ASes.

For this problem, we employ Rocketfuel [26]. The original authors, Mahajan et al., executed `traceroute` probes from about 200 looking glass servers (presented in `www.traceroute.org`) to all the IP prefixes in a chosen AS. Many of the servers are now unavailable; instead, we use 390 `planetlab` [8] nodes, hosted in educational institutions across the globe.

For each chosen AS, we find which prefixes it advertises (on `cidr-report.org`). We then run `traceroute` probes targeted to one representative IP address in every prefix. This gives us a router-level path ending inside the AS.

Next, using Whois [5], we inspect each `traceroute` trace to identify the first IP address belonging to the target AS. These are, potentially, the edge routers of the AS. We trim the traces up to these addresses, as we only need paths inside the AS.

The router IPs (belonging to the target AS), discovered through the above process, are quite noisy and suffer from problems such as anonymity and aliasing [18]. To clean them, we use the state-of-the-art alias resolution tool, Midar [6].

Finally, we combine the paths into a map of the AS, and selected routers appearing in a large number of `traceroute` probes as candidates for firewall replacement.

B.Finding DNS Resolvers

One of the most challenging tasks was to find all the DNS resolvers of the country. Finding DNS Resolvers consists of various steps.First , I collected all the Indian Autonomous Systems.There were total of 884 Autonomous Systems.All the prefixes which these ASes advertise were collected from CIDR report.

Next, Nmap scan was performed in each of these prefixes of each AS on UDP port53(DNS is active on port number 53).Nmap is the network scanner which is used to identify hosts and services on computer network.There were three kinds of results found: open , filtered and closed. Closed corresponds to unreachable message responses from the destination. Open corresponds to those situations where client receives meaningful responses.Filtered corresponds to those responses where client receives no message.

Finally, each IP which was found to be open or filtered on UDP port 53 was used to determine the IP address of popular WWW destinations (e.g. https://www.google.com).If they were able to do resolution , they were added to the list of DNS resolvers.

With this list of DNS Resolvers , using aforementioned approach , we constructed a graph of prefix-to-AS paths connecting the IP prefixes corresponding to DNS Resolvers , and all the Indian ASes.We have also included OpenDNS server and Google DNS server.

Chapter 4

Evaluation and Results

4.1 Experimental Results

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

4.1.1 Network Locations for IP Filtering

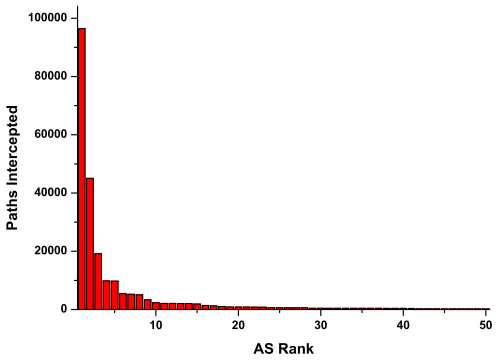

Figure 4.1: Paths intercepted by individual ASes vs AS rank (acc. to paths freq.)

We obtained 186679 paths connecting Indian ASes to the potentially filtered sites. Figure 4.1 represents the number of paths an individual AS intercepts. The horizontal axis of the graph represents ASes ranked according to the number of paths each one intercepts. The ASNs and their owner organizations are presented in the table 4.1. Apparently, some Indian ASes - a very small number - appear in the majority of these paths. The cumulative results of *paths intercepted*

vs *total number of ASes* is presented in figure 4.2.

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 4.1: AS Ranks, their ASNs and their owners.

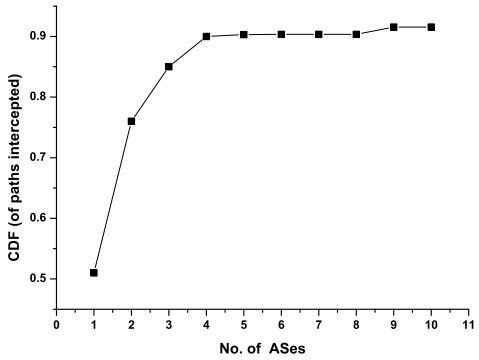

Figure 4.2: CDF of Indian paths intercepted by ASes.

From figure 4.2, *approximately* 4 *ASes can censor over* 90% *of the paths to the censored destinations.*

Identifying key routers inside key ASes: After identifying key ASes in the Internet, our efforts involve identifying the key network routers inside some of the important ASes.It is these routers,which transport a large fraction of network traffic,that we will choose to replace with firewall or dns injector.(our aim is, of course,to use a small number of routers and stll capture all the traffic that originates from or transits through the AS).

As described in the previous section , we run traceroute probes from about 390 planetlab hosts to IP prefixes inside the target ASes.For our experiments, we chose 4 top ASes named- $AS9498$, $AS9583$, $AS10201$, and $AS4755$.

Our traceroute results were cleaned using the midar tool to resolve aliases.From the final results,we obtained statistics regarding the number of routers which appeared most in our traceroute paths.

From our results, in some ASes (AS4755), a few (approx 40) routers appear in over 80% of the traced paths.In others,such as AS9498 , many more routers(> 100) are required to get good path coverage(90% of the paths. Statistics for each AS,including the number of edge and core

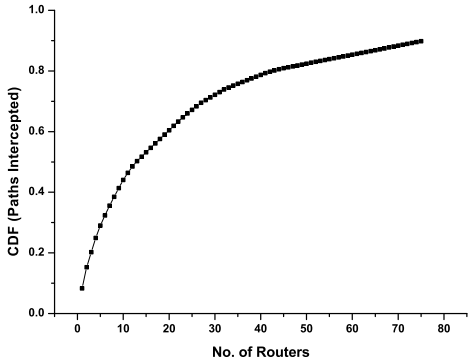

Figure 4.3: CDF of `traceroute` paths intercepted by individual routers, sorted by increasing number of paths through each router(for $AS4755$.)

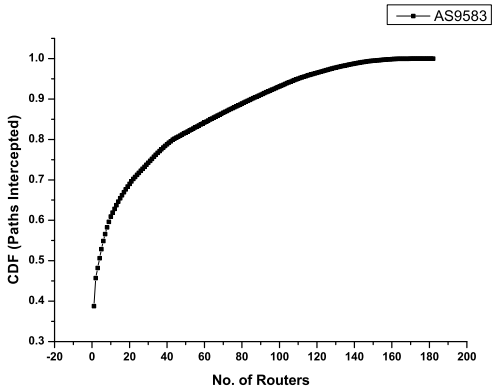

Figure 4.4: CDF of `traceroute` paths intercepted by individual routers, sorted by increasing number of paths through each router(for $AS9583$.)

routers, are represented in 4.2.

The table not only presents the number of edge and core routers,but also those that appear in the most paths,covering a large fraction of the network topology.In case of AS4755,traceroute probes to its 548 advertised prefixes revealed about 319 routers(291 edge + 28 core).We need more than half of edge routers - more precisely (185 edge routers) of the frequently seen routers to cover (most of) the traceroute paths.Similarly , for AS 9498 , we need 269(235 edge + 34 core) out of 519 routers (320 edge + 199 core).

We see that we can do considerably better than the naive solution of replacing the edge routers of the AS. While replacing edge routers would indeed intercept all the traffic entering and leaving an AS,several of the edge routers rarely intercept traffic,while some ("backbone") core routers are much more significant.Our mapping approach allows us to get good coverage with a smaller number of routers.Apparently, about 15 routers cumulatively appear in about only 33% and 52% of the paths corresponding to $AS4755$ and $AS9498$ respectively. At most 5% of the paths transit any router (corresponding to router ranked 1). It implies that for these ASes, no small number of routers can collectively censor a large volume of traffic.

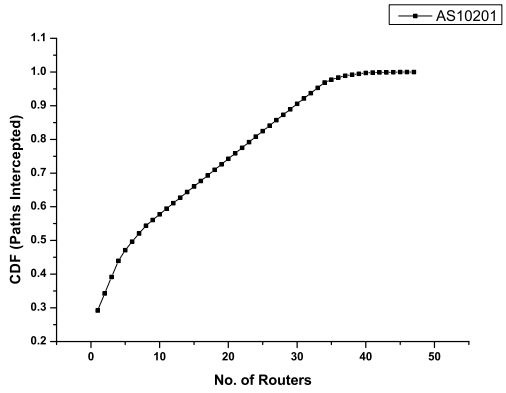

Figure 4.5: CDF of `traceroute` paths intercepted by individual routers, sorted by increasing number of paths through each router(for $AS10201$.)

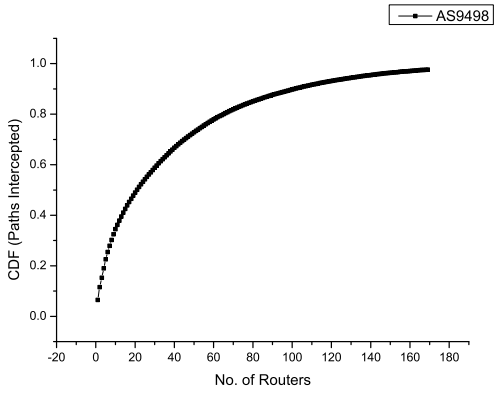

Figure 4.6: CDF of `traceroute` paths intercepted by individual routers, sorted by increasing number of paths through each router(for $AS9498$.)

We used this to create router-level maps of the most significant 4 ASes, $AS9498$ $AS4755$, $AS10201$ and $AS9583$. We sorted the routers by the number of `traceroute` paths they uniquely intercept, and selected those routers that appear on a large number of paths. CDF of paths intercepted are shown in figures 4.3, 4.4, 4.5, 4.6 . (For privacy concerns, we refrain from revealing the IP addresses of these routers.)

ASN	**No. of Edge Routers**	**No. of Core Routers**	**No. of prefixes advertised**
4755	185/291	5/28	548
10201	47/57	0/0	310
9583	182/225	0/13	1113
9498	235/320	34/199	120

Table 4.2: No. of routers (edge and core) inside the AS that appear in all the paths

		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 4.3: IP prefix hijack

4.1.2 Censorship Through IP Prefix Hijack

As described in the previous section, 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, sorted by their node degrees, and determined the number of ASes potentially vulnerable to attacks from each of these ASes. The results of these simulations are presented in table 4.3.

From the table, we see 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, an attack by $AS9498$ can affect a total of 955 ASes (896 Indian and 59 others).

4.1.3 Censorship Through DNS Injection

Using our approach for identifying open DNS resolvers, we identified a total of 55234 publicly accessible DNS servers from probing all IP addresses of India. Detailed process is depicted in figure 4.7.

After identifying the prefixes corresponding to these each resolver IP, we selected one corresponding to each AS[1] Ultimately 355 prefixes, representative of unique 355 Indian ASes, were selected. Using Gao's algorithm, we estimated paths from Indian ASes to prefixes of DNS resolvers in India. Also, paths from Indian Ases to OpenDNS and Google DNS were calculated. *Cumulatively, 8 ASes (according to path frequency) can intercept* 99.14% *of these paths, and potentially launch DNS Injection attacks (see figure 4.8).* These ASes also appear among our top 10 AS choices for IP filtering and IP prefix hijacking.

Collateral Damage: Our graph of paths from censored prefixes to ASes has 186679 paths of Indian origin (1.76% of paths). About 121931 paths of foreign origin (1.15% of paths, *comparable to the fraction originating in India*) transit through or terminate in an Indian AS as shown in figure 4.9. *Censorship by Indian ASes may inadvertently impact a very large number of unintended customers, across Finland, Hong Kong, Singapore, Malaysia, the US etc.*

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

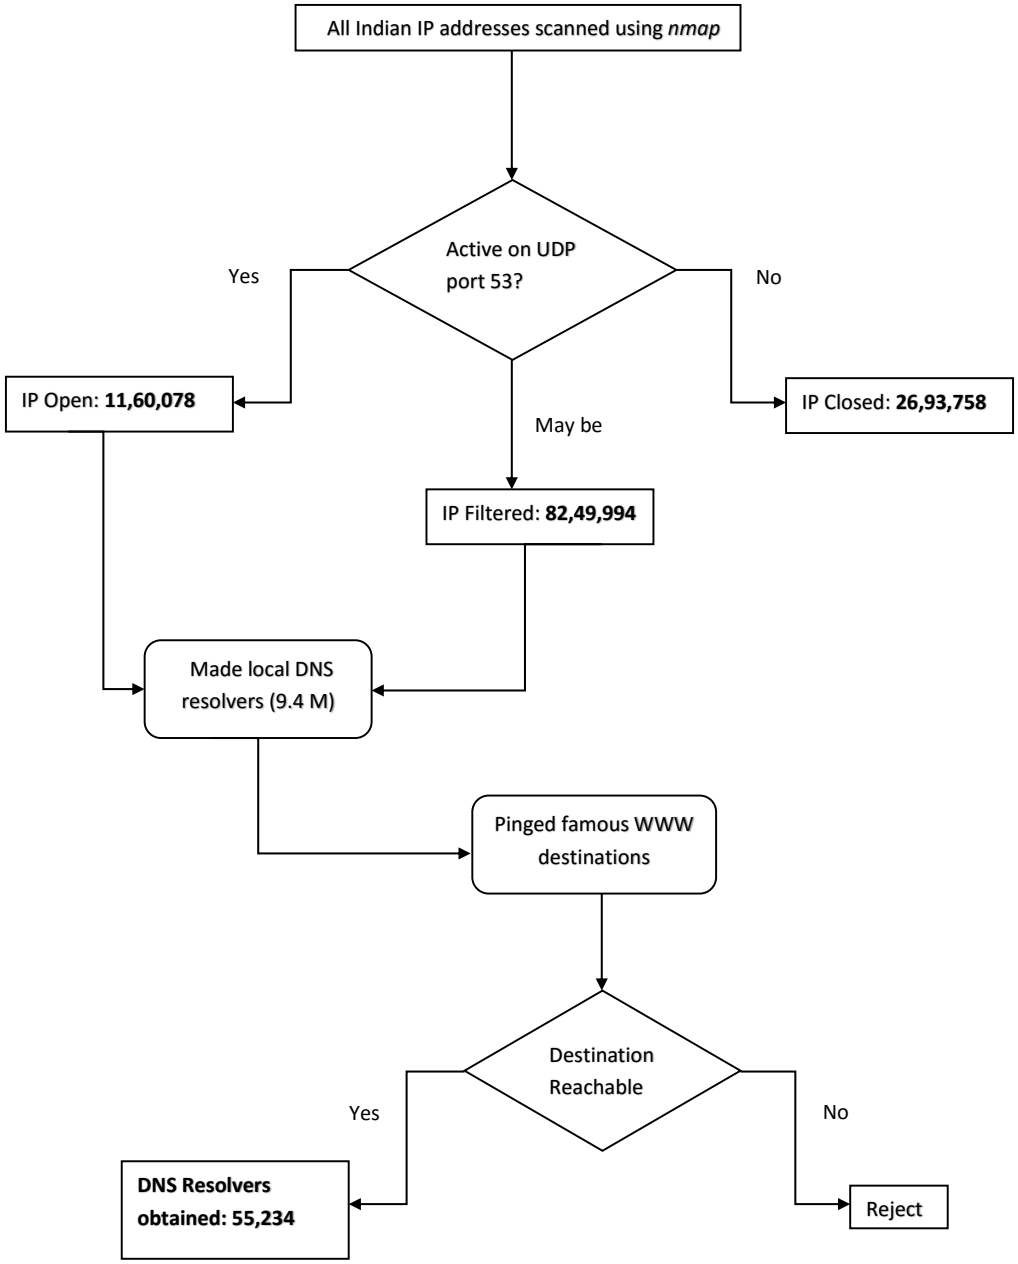

Figure 4.7: Process to find DNS resolvers

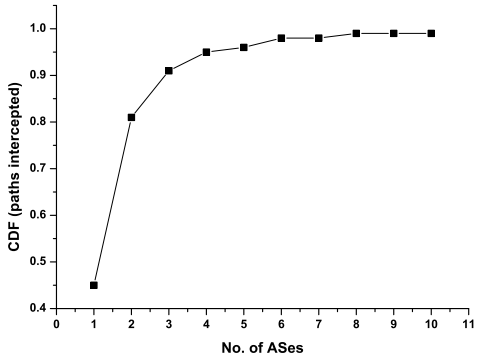

Figure 4.8: CDF of DNS paths intercepted by top 10 Ases.

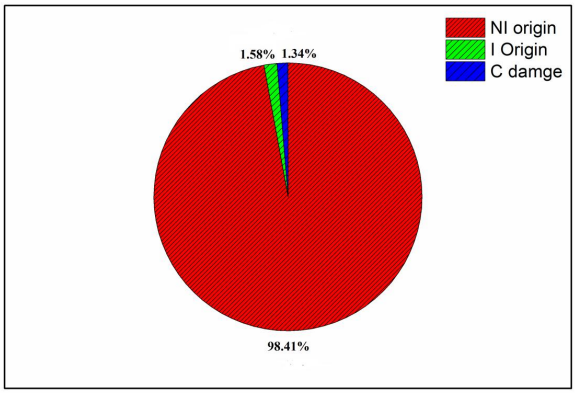

Figure 4.9: Contribution of collateral by India

Chapter 5

Discussion and Future Work

5.1 Discussion

Internet censorship in developing nations is both easier and harder than in the developed world. The Internet infrastructure is smaller, but the government is also resource constrained. In this thesis, we study how this trade-off affects the feasibility of censorship in India, a country which shows signs of becoming censorious [4].

Previous work on censorship in poorer countries [13, 15, 21] focuses on the content (*what* is being censored), and the mechanism (*how* censorship is implemented). As large-scale censorship is not yet in effect in India, we have taken a list of likely future targets, and focused on the mechanisms of IP filtering, DNS injection, and IP prefix hijacking. We find that there are indeed a few key ASes (less than 10) where an adversary with limited budget might deploy its infrastructure to devastating effect. Another surprising result is the extent of collateral damage - many foreign ASes can be affected by Indian censorship, including a substantial portion of Africa, and several "free" countries such as Norway and the US.

5.2 Future Work

It is important to note that we make use of no features peculiar to India. Our analysis may, therefore, be applied for other countries as well. In our immediate future work, we intend to do this; our long-term goal is to develop metrics for how hard it is to censor traffic in a country, as well as how "central" a country is, i.e. how much collateral damage it can cause.

Our analysis may be extended in several ways. For example, what happens if a country blocks different content - search engines and social networking sites (as seen in China), instead of pornography (or other objectionable content)? Also, in practice, targeted content is often hosted on social media sites, or other sites with apparently benign URLs; a real censor does not just block IP addresses, it performs real-time pattern recognition w.r.t. content. (Semantics-based filtering is very hard: attempts to block jihadi sites also block sites that monitor militancy, such

as `jihadwatch.org`.) How might a sophisticated but resource-poor adversary perform such censorship? We intend to study these questions in our future work.

Chapter 6

Conclusion

6.1 Concluding Remarks.

In this thesis, we have applied a novel method of analysis to study censorship (by a resource-restricted adversary) as an engineering problem. Our results show that it is indeed feasible for the Indian Government to become a censorious regime. *However, while it is sufficient for a censor to assume control over a small fraction of Indian ASes, it is not enough to control a small proportion of the routers in the ASes.*

1. *Potential locations for monitoring and filtering network traffic:* The adversary, given control over the routers of only 4 ASes, can observe 90.52% of the paths to our sample censored sites.

2. *Potential ASes to filter DNS requests:* With control of 8 ASes, the adversary can deny access to alternative DNS servers for 99% of the users.

3. *IP Prefix Hijacking:* There exist several ASes which can single-handedly launch IP Prefix hijacking attacks that affect all Indian users.

Further, there are ASes outside India whose traffic may be inadvertently censored if the Indian government employs these means of censorship.

Bibliography

[1] Alexa - Actionable Analytics for the Web. http://www.alexa.com/.

[2] Archipelago (ark) measurement infrastructure. http://www.caida.org/projects/ark/.

[3] Herdict:Help Spot Web Blockages. http://herdict.org/.

[4] Internet Censorship in India. https://en.wikipedia.org/wiki/Internet_censorship_in_India.

[5] Ip to asn mapping. http://www.team-cymru.org/IP-ASN-mapping.html.

[6] Midar. http://www.caida.org/tools/measurement/midar/.

[7] Pakistan hijacks youtube. http://research.dyn.com/2008/02/pakistan-hijacks-youtube-1/.

[8] Planetlab – an open platform for developing, deploying and accessing planetary-scale services. https://www.planet-lab.org/.

[9] Porn websites blocked in india: Government plans ombudsman for online content. http://gadgets.ndtv.com/internet/news/porn-websites-blocked-in-india-government-plans-\ombudsman-for-online-content-723485.

[10] Route views project. http://archive.routeviews.org/.

[11] Traceroute Looking Glass.

[12] University of Oregon Route Views Project. http://www.routeviews.org/, 2000.

[13] 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, 2012), USENIX.

[14] Anonymous. The collateral damage of internet censorship by dns injection. *SIGCOMM Comput. Commun. Rev. 42*, 3 (June 2012), 21–27.

[15] Aryan, S., Aryan, H., and Halderman, J. A. 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, 2013), USENIX.

[16] Ballani, H., Francis, P., and Zhang, X. A study of prefix hijacking and interception in the internet. *SIGCOMM Comput. Commun. Rev. 37*, 4 (Aug. 2007), 265–276.

[17] Gao, L. On inferring autonomous system relationships in the internet. *IEEE/ACM Trans. Netw. 9*, 6 (Dec. 2001), 733–745.

[18] Haddadi, H., Rio, M., Iannaccone, G., Moore, A., and Mortier, R. Network topologies: inference, modeling, and generation. *IEEE Communications Surveys Tutorials 10*, 2 (Second 2008), 48–69.

[19] Levis, P. The collateral damage of internet censorship by dns injection. *ACM SIGCOMM CCR 42*, 3 (2012).

[20] Lyon, G. Nmap: the network mapper - free security scanner. http://insecure.org/fyodor/.

[21] Nabi, Z. 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, 2013), USENIX.

[22] (NSA), U. S. N. S. A. Prism (surveillance program). https://en.wikipedia.org/wiki/PRISM_(surveillance_program).

[23] Qiu, J., and Gao, L. As path inference by exploiting known as paths. Tech. rep., In Proceedings of IEEE GLOBECOM, 2005.

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

[25] Rocketfuel: An ISP Topology Mapping Engine. http://www.cs.washington.edu/research/networking/rocketfuel/.

[26] Spring, N., Mahajan, R., Wetherall, D., and Anderson, T. Measuring isp topologies with rocketfuel. *IEEE/ACM Trans. Netw. 12*, 1 (Feb. 2004), 2–16.

[27] Winter, P., and Lindskog, S. How the Great Firewall of China is blocking Tor. In *Proceedings of the USENIX Workshop on Free and Open Communications on the Internet (FOCI 2012)* (August 2012).