TASVEER : Tomography of India's Internet Infrastructure

Student Name: Rahul Kumar Singh

IIIT-D-MTech-CS-GEN-MT-13-101
Jul 30, 2015

Indraprastha Institute of Information Technology
New Delhi

Thesis Committee
Sambuddho Chakravarty (Advisor)
Vinayak Naik (Internal Examiner)
Mohan Dhawan (External Examiner)

Submitted in partial fulfillment of the requirements
for the Degree of M.Tech. in Computer Science,
with specialization in Information Security

Keywords: autonomous system, AS level path, router level path, traceroute, centrality, BGP routing and network topology

Certificate

This is to certify that the thesis titled **"Tasveer : Tomography of India's Internet Infrastructure"** submitted by **Rahul Kumar Singh** 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 him under my guidance and supervision in the CERC group at Indraprastha Institute of Information Technology, Delhi. This work has not been submitted anywhere else for the reward of any other degree.

Professor Sambuddho Chakravarty
Indraprastha Institute of Information Technology, New Delhi

Abstract

With approximately 250 million Internet users, India stands amongst the top 5 Internet using nations of the world. India's network space is made up of 789 Autonomous Systems (ASes), that route all the network traffic of India. On the other hand, US has approximately 300 million users, whose traffic is routed over 22K ASes. Thus, a relatively small network routes the traffic of large number of Indian users. Failures and attacks in such networks could impact large number of users. However, being a relatively small number, it becomes easy to generate maps presenting the connectivity of ASes in the networks and the routers that make up the ASes. Such information could be used for various purposes such as diagnosing network failures and attacks, large scale network surveillance and bypassing such surveillance, load balancing, efficient content distribution and delivery.

We present, a first effort to our knowledge, the topological information of India's entire Internet space representing the connectivity between all 789 ASes and intra-domain routers. Our research presents information of routers and ASes that transport relatively large fraction of traffic for vital network installations like popular ISP users, important organizations like financial institutions, educational institutions, research organizations etc.

Acknowledgments

I would like to thank Dr. Sambuddho Chakravarty for his unparalleled guidance and motivation. Without his guidance this work would not have been possible. The amount of time that he has taken out from his busy schedule is quite commendable and worth mentioning. His knowledge on the topic and his motivation and patience to lead us towards quality work is overwhelming source of inspiration. It's an honor and a privilege to have him as my mentor. I thank him from the bottom of my heart.

My heartfelt thanks to Dr. Mohan Dhawan and all the members of CERC(Cybersecurity Education and Research Centre) for their valuable suggestions. A special thanks to all my friend for critically reviewing and commenting on the work and helping to make it what it is. I would like to thank my parents for always supporting me and providing me an environment where I could work dedicatedly. Last but not the least, I would like to thank IIIT Delhi 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

The organization of the networks in the developing regions of the world has been a subject of study in recent past. Several recent research efforts have tried to look into the problem of characterization of network traffic and topology of networks for developing nations [13] . However few have tried to look into a full scale network topology measurements.

Through our research we tried to study the topology of networks in India i.e the organization of Autonomous Systems (ASes) and routers in India. The overall questions we tried to answer are such as: *Is there some small set of ASes and routers through that transports a large fraction of India's Internet traffic ?*

Gathering information about the Internet topology is a mandatory task for modelling the network but also for monitoring the network. Moreover, several network applications, such as *network diagnosis, server selection, overlay routing, content distribution* etc can benefit from the knowledge of topology of networks.

Today, an Internet can be so large that one routing protocol cannot handle the task of updating the routing table of all routers. For this reason , an Internet is divided into autonomous systems. An autonomous systems(AS) is a group of networks and routers under the authority of a single administration. Routing inside an autonomous system is referred to as intra domain routing. Routing between autonomous systems is referred to as inter-domain routing. Each autonomous systems can choose one or more intra domain routing protocols to handle routing inside the autonomous systems. However, only one interdomain routing protocols handles routing between autonomous systems.

Prior work [9] broadly classifies ASes into three categories, namely customers, providers and peers. ASes that typically derive network connectivity from other ASes are customers, e.g., individuals and organizations that are typically sources and sinks of network connections. Individual Internet users, banks, several small and large commercial and non-commercial organizations etc., rely on other ASes to connect to the Internet. However, these ASes do not generally relay network traffic for other ASes (like ISPs), and are typically classified as providers. In contrast, ISPs could be classified as both providers and customers simultaneously. Finally, ASes may also share peer-to-peer relationships. These ASes are no different from ordinary ASes, but unlike

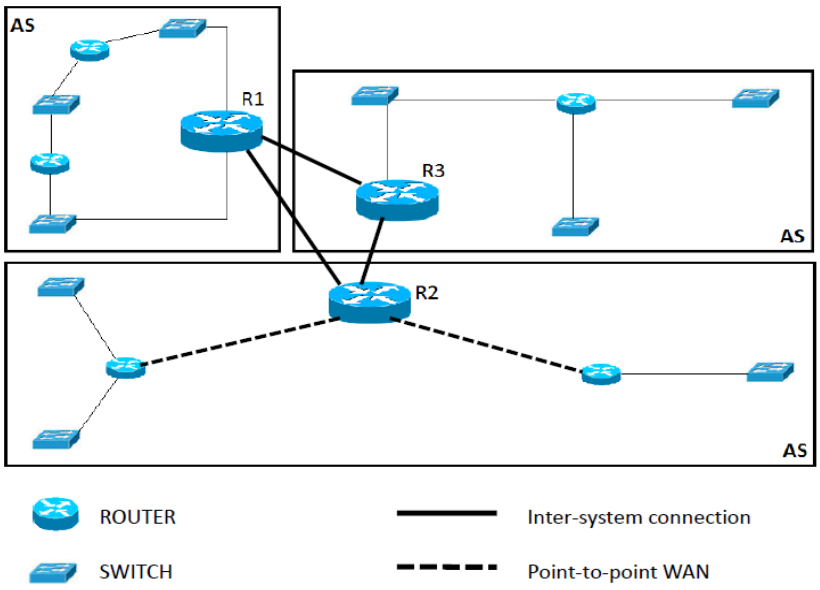

Figure 1.1: Inter and Intra domain routing(Autonomous Systems)

the customer-provider relationships, these ASes mutually agree upon relaying each others traffic without any monetary incentive.Prior work [9] also defines valid or invalid paths based on economic dependence of ASes on one another. A valid path between source and destination ASes is where for every ISP providing transit, there exists a payee AS and the payee of the transit provider must be its immediate neighbor in the path The AS path estimation mechanism [15] makes use of such relationships.

India which is the second largest nation by population has approximately 250 million Internet users with 789 ASes whereas US has over 22k ASes to route traffic for their 300 million Internet users. Hence, a very small network infrastructure is used in India to route the traffic of their Internet users.Thus, we are trying to find out that with such a relatively small network infrastructure does there exists some nodes(AS and router) in the network which is transporting a major or large fraction of Indian Internet traffic.

List of Contributions : our work makes the following three contributions to tomography of India's Internet infrastructure:-

- we develop the AS connectivity graph of India and identify few critical ASes based on that graph.
- we identify all the critical routers of top three ASes of India and infer some other interesting statistics.

- we identify all the critical routers which transport major fraction of traffic to important network installations in the country such as financial institutions, defence websites, banks, e-commerce websites etc.

1.1 Problem and Approach

The goal of our work is to obtain AS- and router-level topology of an entire nation such as India. In this section, we describe what we mean by an AS and router level topology and our approach to obtain the targeted topology. Internet is divided into autonomous systems. An autonomous systems(AS)is a group of networks and routers under the authority of a single administration. Maps showing connectivity between these ASes are known as AS-level connectivity graph and maps showing connectivity between routers inside these ASes are known as router-level connectivity graph. Our aim is to discover AS and router connectivity maps and infer critical information from that maps such as - *does there exists small set of AS and routers which transports major fraction of Internet traffic.* Apart from this, a major concern is related to network and communication security. An attack on routers and ASes which transports major fraction of Internet traffic could be catastrophic. An adversary could bring down a major part of network by targeting these critical routers/ASes or an adversary could merely snoop on several network flows by observing flows entering and leaving only these routers. Thus, divination of such critical routers/ASes is must. Generating AS- and router-level connectivity maps for an entire nation opens up gate for several other security questions which we will discuss later.

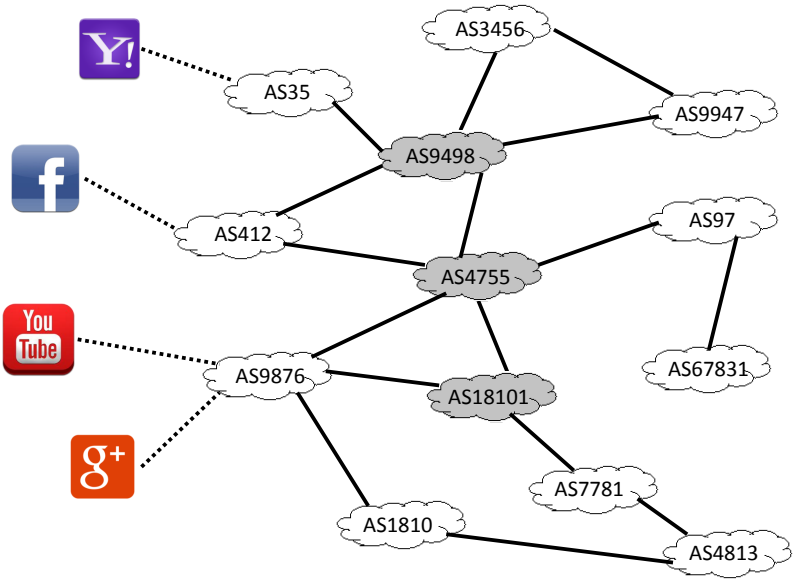

Figure 1.2: AS connectivity graph with AS9498, AS4755 and AS18101 are the critical ASes as most of the traffic passes through these ASes

For AS level connectivity maps we are using Gao *et al* [15] MULTIPATHS algorithm which relies on the passive analysis of BGP routing information. For router level connectivity, we are taking leverage from Rocketfuel paper [17] which uses publicly available traceroute servers to launch traceroute probes to map pop-level topology. As most of the traceroute servers are unavailable, therefore we are using ∼500 PlanetLab [4] nodes to launch traceroute from various locations across the globe to map routers of the targeted network.

Even after the connectivity information has been obtained through Gao *et al* [15] and traceroutes, two difficulties remain. First, each traceroute consists of a list of IP addresses that represent router interfaces. For an accurate map, the IP addresses that belong to the same router, called *aliases*, must be resolved. For resolving aliases, we are using MIDAR tool [12].

Second, to find the critical AS/router which are critical or through which major fraction of entire traffic flows through, we are using peer centrality algorithm [2] which will run on the graph obtained from our mapping efforts.

1.2 Background

Generating maps of the Internet, showing connectivity of routers and ASes has been an active area of research for the past decade and a half. Donnet *et.al.* [7] provide a comprehensive survey of various mapping efforts, classifying them into categories based on different topological abstractions, such as *link layer, network layer, Internet* and *overlay*.

Unlike other abstractions, network layer topology specifically identifies connectivity between network layer devices such as routers, primarily utilizing tools like traceroute to build maps of AS- and router-level connectivity. However such network level tomography is susceptible to errors due to traceroute itself. Note that routers have multiple network interfaces, often associated with different network subnets . Thus traceroute probes may list these interfaces as belonging to different routers(since at each interface the TTL of IP packets is decremented), thereby polluting the tomography. These different interfaces of the same router are known as *aliases*, and precise network tomography requires such router level alias resolution [11, 12].

1.2.1 Autonomous Systems

An Autonomous System (AS) is a collection of routers whose prefixes and routing policies are under common administrative control. This could be a network service provider, a large company, a university, a division of a company, or a group of companies. The AS represents a connected group of one or more blocks of IP addresses (called IP prefixes) that have been assigned to that organization and provides a single routing policy to systems outside the AS. Autonomous Systems create a two-level hierarchy for routing in the Internet. Routing between Autonomous Systems (inter-AS routing) is external to the AS allows one AS to send traffic to another AS. Note that most organizations do not interconnect via autonomous systems but

simply connect via a single ISP. Routers within an AS use an Interior Gateway Protocol (IGP), which handles routing between nodes inside the AS. Common interior gateway protocols include RIP, OSPF, IS-IS, EIGRP, as well as some proprietary protocols such as IGRP. Routing within an Autonomous System (intra-AS routing) is internal to that AS and invisible to those outside it. The AS administrator decides what routing algorithm should run within it. To get traffic from a host in one AS to a host in another AS, the autonomous systems need to be connected. Most ASes do not share a direct link with each other, in which case data traffic may be routed through the networks of other ASes that agree to carry the traffic. An Exterior Gateway Protocol (EGP) is a routing protocol that handles routing between Autonomous Systems (inter-AS routing).

1.2.2 Internet Tiers, Peering, and Transit

Since most ASes are not connected with each other, they need to route their traffic through other ASes.

Peering is when a pair of ASes establish a reciprocal agreement to connect with each other to exchange traffic with each other, without charge. The assumption that each has an interest in connecting to the others customers, similar to how postal systems throughout the world do not charge when mail is routed from one country to another.

Tier 1 ISPs are those that do not have to pay any other network for transit. They peer with all other tier 1 networks (there are only about fourteen of these worldwide). Given any IP address, a Tier 1 ISP will be able to connect directly to a top-level ISP that can route to that address. The United States has eight interconnection regions[2] that create a default free zone where Tier 1 ISPs connect their networks together in peering relationships.

Peering agreements are not necessarily transitive. If AS1 peers with AS2 and AS2 peers with AS3, AS2 is not necessarily obligated to carry traffic to AS3. Whether this is permitted or not is a business, rather than technical, decision.

A Transit relationship is when an ISP (an AS) sells access to the Internet. It is when an AS agrees to act as a router, carrying traffic from one AS and out to some other AS to which it has a link. The complete data path may, of course involve multiple transit hops through different ASes. An AS will typically meter the traffic on each link and charge a transit fee. Depending on policy, an organization in one AS may be charged for traffic even to the connected AS.

A Tier 2 ISP is one that needs to purchase Transit to connect to at least some part of the Internet. Because of transit fees, many Tier 2 ISPs will try to establish peering relations directly with as many Tier 1 and other Tier 2 ISPs as they can so they can exchange traffic with those ISPs for no fee (although there may be a peering fee). For example, it is common for cable companies to peer with content providers such as Google, Amazon, and Microsoft.

1.2.3 AS Relationships

BGP routing relies on inter AS relationship, the business relationships between ASes can be broadly classified into three main categories: *provider-to-customer(p2c), customer-to-provider(c2p)* and *peer-to-peer(p2p)*. In Figure 1.3 the direction of arrow indicates cash flows between ASes. ASes at lower levels pay ISPs at higher levels in exchange for access to the rest of the Internet. The customer ISP pays the provider ISP for transit. A p2p link connects two ISPs who have agreed to exchange traffic on a quid pro quo basis. Peers exchange traffic only between each other and each other's customers. This relationship allows peering ISPs to save money on transit costs.

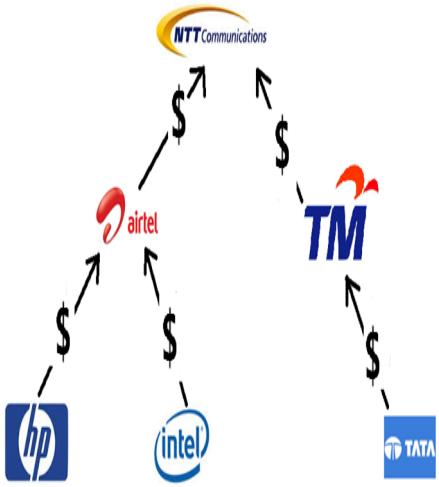

Figure 1.3: AS relationship

1.2.4 Valley Free Path(VALID PATH)

According to Gao *et al* [15] a valid path between source and destination ASes is one in which for every ISP providing transit (a transit provider), there is a payee. The payee of the transit provider must be its immediate neighbor.

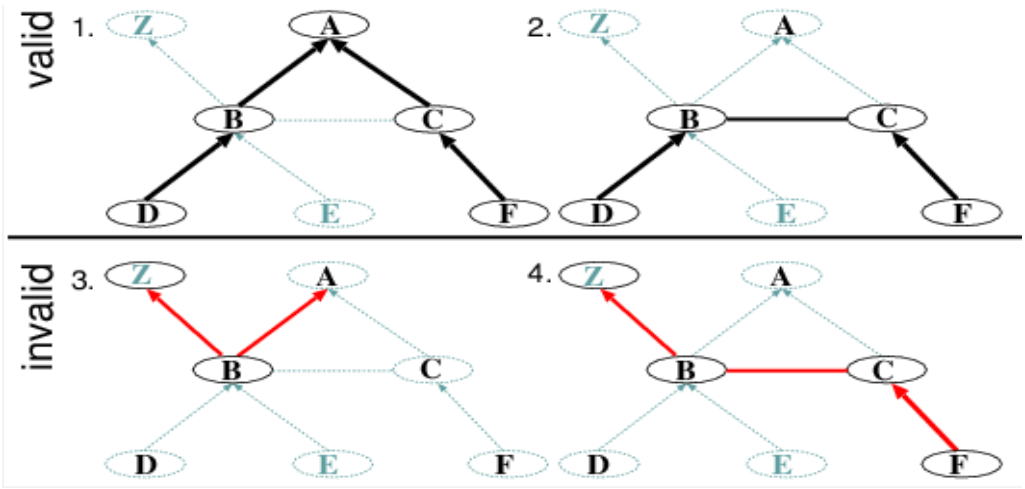

Figure 1.4: Valid and Invalid paths w.r.t valley free property

In Figure 1.4 the top two examples are valid paths, while the bottom two are invalid. In Example 1 the transit providers are A, B, and C. ISPs B and C pay to A, D pays to B, and F pays to C. In Example 2 the transit providers are B and C, and they are paid by D and F respectively. In contrast, in Example 3 the transit provider is B, but not only does no one pay B, but B itself pays both A and Z. Example 4 also illustrates situation where nobody pays transit provider B.

Hence a valid path must have the following valid path pattern: zero or more c2p links, followed by zero or one p2p link, followed by zero or more p2c links.

1.2.5 Traceroute

When computers communicate over the Internet, there are often many connections made along the way. This is because the Internet is made up of a network of networks, and two different computers may be on two separate networks in different parts of the world. Therefore, if a computer is to communicate with another system on the Internet, it must send data through a series of small networks, eventually getting to the Internet backbone, and then again traveling to a smaller network where the destination computer resides.

Traceroute is a TCP/IP utility that allows a user to trace a network connection from one location to another, recording every hop along the way. When a traceroute is run, it returns a list of network hops and displays the host name and IP address of each connection. It also returns the amount of time it took for each connection to take place (usually in milliseconds). This shows if there were any delays in establishing the connection. Therefore, if a network connection is slow or unresponsive, a traceroute can often explain why the problem exists and also show the location of the problem. Figure 1.5 shows the result of traceroute to facebook.com website.

traceroute to facebook.com (173.252.120.6), 30 hops max, 28 byte packets
1 node2-1.realcloud.nl (62.221.205.11) 0.036 ms 0.015 ms 0.012 ms
2 hosted-by.2is.nl (62.221.254.1) 0.205 ms 0.196 ms 0.185 ms
3 2is-10g.nl.jointtransit.nl (217.170.20.187) 0.866 ms 0.861 ms 0.852 ms
4 nikhef-ixr.openpeering.nl (217.170.0.242) 0.874 ms 0.870 ms 0.860 ms
5 er1.ams1.nl.above.net (80.249.208.122) 1.450 ms 1.418 ms 1.416 ms
6 ae14.cr1.ams10.nl.zip.zayo.com (64.125.21.77) 1.886 ms 1.978 ms 1.680 ms
7 v122.ae29.cr2.lga5.us.zip.zayo.com (64.125.30.168) 72.071 ms 72.449 ms 72.439 ms
8 ae1.cr1.lga5.us.zip.zayo.com (64.125.29.37) 72.395 ms 72.328 ms 72.291 ms
9 ae2.er3.lga5.us.zip.zayo.com (64.125.31.214) 72.414 ms 72.156 ms 72.236 ms
10 IPYX-100687-870-ZYO.zip.zayo.com (128.177.165.234) 74.356 ms 74.344 ms 74.278 ms
11 be1.bb01.lga1.tfbnw.net (204.15.20.190) 94.558 ms 95.122 ms 95.180 ms
12 be29.bb01.dca1.tfbnw.net (74.119.78.78) 94.991 ms 95.247 ms 94.861 ms
13 be42.bb01.frc3.tfbnw.net (31.13.26.251) 95.569 ms 95.539 ms 94.751 ms
14 ae62.dr02.frc3.tfbnw.net (173.252.65.103) 94.197 ms 94.911 ms 94.505 ms
15 edge-star-shv-12-frc3.facebook.com (173.252.120.6) 94.320 ms 94.504 ms 94.251 ms

Figure 1.5: Snapshot of traceroute

1.3 Motivation

India has approximately 250 million users today making it the second largest Internet market. A key goal of our research is to understand the network topology and router interconnection for Internet traffic in India. Specifically, our research aims to identify topological properties and anomalies observed for different ISPs at both AS- and router-level. For example, network operators would greatly benefit from identifying which ASes/routers relay the network traffic to large fraction of Indias Internet users, and are thus required to be available all the time etc. Such information could also be useful for purposes of fault isolation and assisted (or automated) failover.

We analyzed the Hurricane Electric database [3] and observed that even in a large developing nation like India, with dense penetration of the Internet, there are few just about 789 ASes. Amongst these ASes there are about 708 customer only ASes. The rest 76 ASes are both transit and origin ASes. These ASes include about 49 ISPs, with 13 of them having providers outside India. These ASes connect India's Internet traffic to the world outside. Thus, there are potentially very few ASes that transport large fraction of India's Internet traffic.

Majority of the ASes in India are customers home Ases, commercial and private institutions etc. There have been very few efforts to study the topological connectivity of Ases within India and the connectivity of routers within developing nations like India.

Some network mapping efforts, such as Rocketfuel [17] have tried to study the PoP-level topology of ISPs, in contrast our research focuses on whole tomography (inter and intra domain connectivity) of developing nations such as India and involves finding out the important/critical network locations: Autonomous systems and routers that carry considerably large fraction of network traffic.

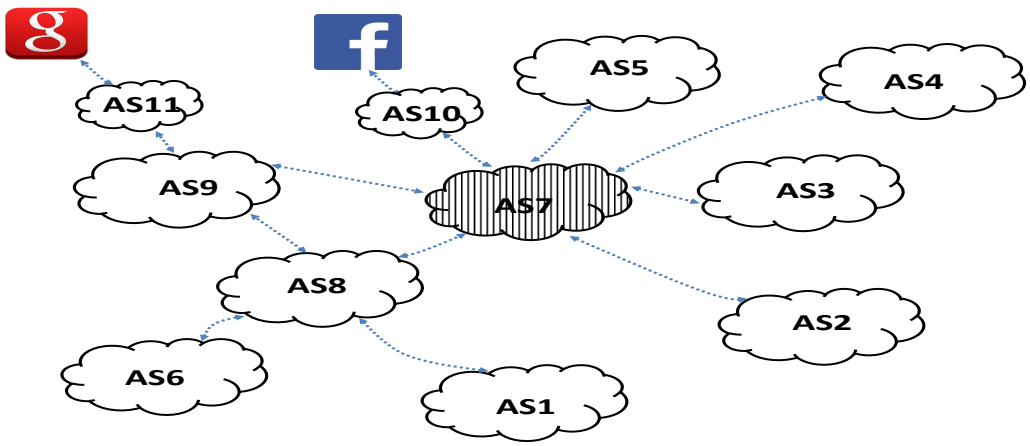

Figure 1.6: Monitor small set of networks/Autonomous Systems/IXes/routers to gather maximum possible information.

However, after over a decade, the network providers and key players in the business have changed. There are several new ISPs and network connections. Thus we believe it is very essential to do a large scale study of the network connections within India.

Further, a larger concern is related to network and communication security. Potentially a small fraction of network routers and Ases transport large fraction of the entire network traffic. Thus an attack on these routers and ASes could bring down the entire network. Alternatively, an adversary could merely snoop on several network flows by observing flows entering and leaving only these routers. Thus identification of such critical routers and ASes that potentially carry major fraction of network traffic is of prime importance.

Chapter 2

Related Work

While network topology mapping is a well studied area of research,there have been few large scale efforts to map large nationwide networks. The network mapping can be broadly classified into two categories: AS-level topology and router-level topology.

Based on traceroute, a number of techniques and tools have been developed to discover the network topology. While the Rocketfuel tool [17] which measures the pop-level ISP topologies has been available for over a decade now, there have been few efforts to use such tools to map large scale networks (particularly in developing nations like India).

Skitter [14], developed by CAIDA, provides a relatively complete view of the Internet AS-level topology. Skitter records paths from a source to many destinations using parallel ICMP traceroutes and stores the replies from each router on the path to the destination host, along with the round-trip times(RTTs).

Mercator [10] discovers the Internet router-level topologies. It applies novel alias resolution heuristics and takes advantage of source-route capable routers to enhance the accuracy of the map. Mao et al developed tools to discover AS-level forwarding with some systematic IP-to-AS mapping techniques.

Branigan *et al* [6] use a simple traceroute to map the internet. They randomly select a target host from every network announced via BGP to the internet core or listed in various internet databases. They stop discovering a route when it contains at least two unresponsive hops or when a firewall is encountered. They scan about 10% of the list of networks each day and scan the entire list on the first of each month.

The Distributed Internet Measurement and Simulations(DIMES) [16] system is a measurement infrastructure that achieves a large scale by following the model of SETI@home. DIMES provides a publicly downloadable route tracing tool, with incentives for users. The DIMES agent performs internet measurements such as traceroute and ping at a low rate, consuming at peak 1KB/sec.

Scriptroute [18] is a system that allows an ordinary internet user to perform network measurements from several distributed vantage points. It proposes remote measurements execution of

PlanetLab [4] nodes through a demon that implements ping, traceroute, hop-by-hop bandwidth measurement, and a number of other utilities.

Internet topology inference based on BGP routing information is another approach. BGP routing tables are widely exploited to construct Internet AS-level topologies. Gao *et al* [15] AS mapping methods which is heavily rely on the passive analysis of BGP routing information have been applied in by Anonymity and Privacy community. In those efforts, the authors used Gao *et al* [15] AS mapping methods for mapping the ASes that appear in paths leading to Tor entry relays and leading out of exit relays. Faloutsos *et al* [8] unveiled Power-Laws in the internet topologies.

Chapter 3

Approach for Gathering Data

We gather data for tomography of India's network space in two phases. First, we generate AS connectivity maps that represent the ASes that most network traffic originating from India would traverse while communicating to popular Internet destinations. Specifically, we leverage the Lixin Gao MULTIPATHS algorithm [15] to generate paths from a known IP prefix to every AS in India. We used IP prefixes of all the popular network destinations of Indian Internet users, as available in the Alexa database [1]. Second, we generate connectivity maps of all routers within each AS using the Rocketfuel tool [17]. Rocketfuel [17] relies on traceroute from looking glass servers across the globe. However, as of April 2015, a majority of the 200 plus looking glass severs were not available during the time of our experiments. We thus launched the traceroute probes from ∼500 PlanetLab [4] machines distributed across the globe.

3.1 Inter-AS path estimation(AS connectivity)

As mentioned previously, we analyzed the publicly available Hurricane Electric BGP database [3] and observed that even for a large developing nation like India, with dense penetration of the Internet, there are only 789 ASes carrying India's Internet traffic to the outside world. Amongst these ASes there are about 708 customer ASes, while there are just about 49 ISPs (includes large and small ISPs, data and voice providers etc.). We selected the top 500 most visited websites (from India as presented by the Alexa database [1]) to infer the entire paths for all Indian ASes from those prefixes. These paths either terminated in India or transited via Indian ASes to other downstream customer ASes. We leverage the MULTIPATHS algorithm [15] to estimate the path between an IP prefix and every other AS. The algorithm takes existing paths to each prefix and extends those using a depth first traversal to discover paths to various ASes, starting at ASes obtained from BGP Routing Information on Bases (RIBs) gathered from Routeviews datasets [5]. The Routeviews datasets present RIBs gathered from Internet Exchange Points (IXes) where several networks (mostly large ISPs and ASes) peer.

Paths corresponding to the RIBs are known as *sure paths*, and the ASes along these sure paths

1.0.4.0/24 1239 174 7545 56203
1.0.4.0/24 3257 174 7545 56203
1.0.4.0/24 293 2828 7545 56203
1.0.4.0/24 1668 6453 7545 56203
1.0.4.0/24 6762 174 7545 56203
1.0.4.0/24 22652 6939 7545 56203
1.0.4.0/24 37100 6939 7545 56203
1.0.4.0/24 2914 2828 7545 56203
1.0.4.0/24 286 6939 7545 56203
1.0.4.0/24 3549 6939 6939 7545 56203
1.0.4.0/24 3741 6939 7545 56203
1.0.4.0/24 6539 577 2828 7545 56203
1.0.4.0/24 1221 4826 38803 56203
1.0.4.0/24 6939 7545 56203
1.0.4.0/24 852 7545 56203
1.0.4.0/24 2152 2828 7545 56203
1.0.4.0/24 5413 2828 7545 56203
1.0.26.0/23 1239 2914 2519

Figure 3.1: Snapshot of RIBs

are known as *Base ASes*. The algorithm extends these sure paths to other ASes (provided the extended path satisfies the valley-free property) to which there are no explicitly known paths (from a given prefix). For example, consider the known route from RIBs $R_k = v_k$,v_{k-1} ... v_{1p}, where v_i is an AS and p is a destination prefix and v_{1p} is the AS that has the prefix p. Based on the propagation of route updates via BGP, we know that R_k is derived from the best path $R_{k-1} = v_{k-1}$, ... v_{1p} from AS v_{k-1} to prefix p. Thus, from R_k, we can extract *(k-1)* sure paths from a known path of length k.

Once the AS-prefix graph has been generated (from the Routeviews dataset) and all the sure paths combined, the graph can be further extended only if the extended path is both valley-free and loop-free. For example, in Figure 3.2 the sure path from AS30 to AS9812 can be extended to also include AS88, and further AS5612. The algorithm uses monotonic path ranks to select the paths to be extended. The path selection criteria sorts the paths in decreasing order of preference based on (i) total path length, (ii) extended path length, and (iii) how frequently the sure path (that is being extended) is observed in Routeviews datasets, and finally returns the best one.

Once the paths from the chosen prefixes to the Indian ASes have been inferred, it results in a graph wherein the vertices are the ASes and the edges are the connections between these ASes(based on the inter AS relationships). We then applied peer centrality [2] on this AS connectivity graph to determine the central(major) ASes or in other words ASes that potentially carry major fraction of India's Internet traffic.

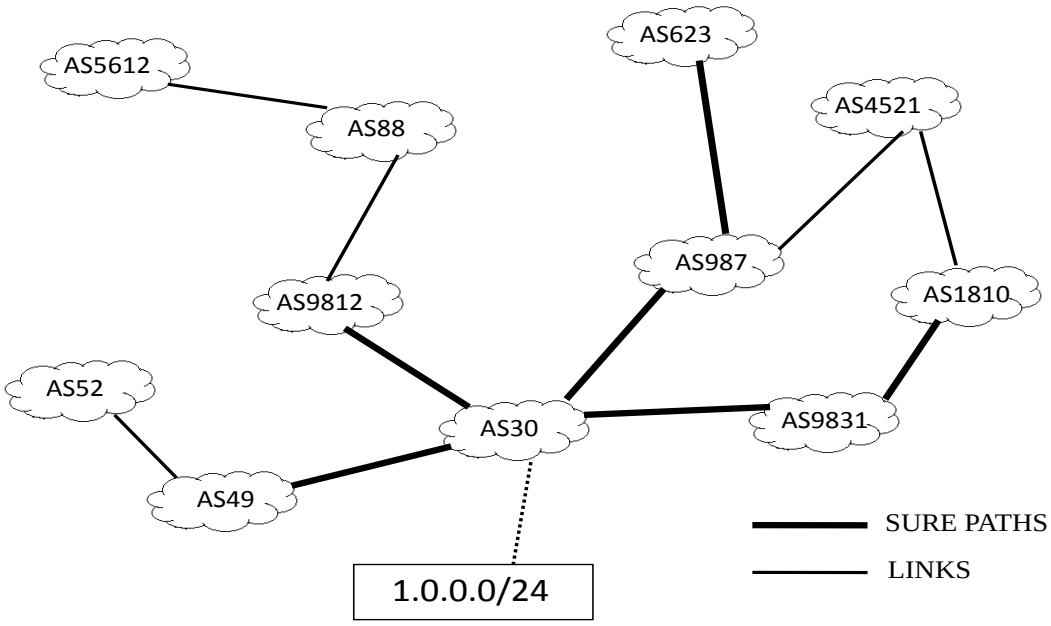

Figure 3.2: Identifying IP prefix to AS path

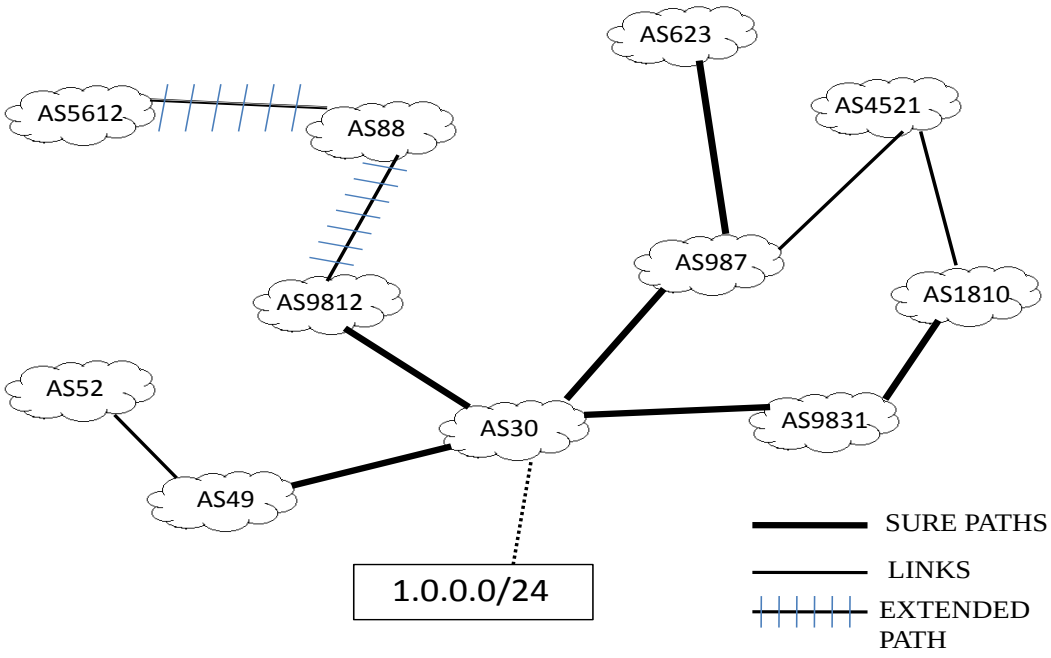

Figure 3.3: Identifying IP prefix to AS path

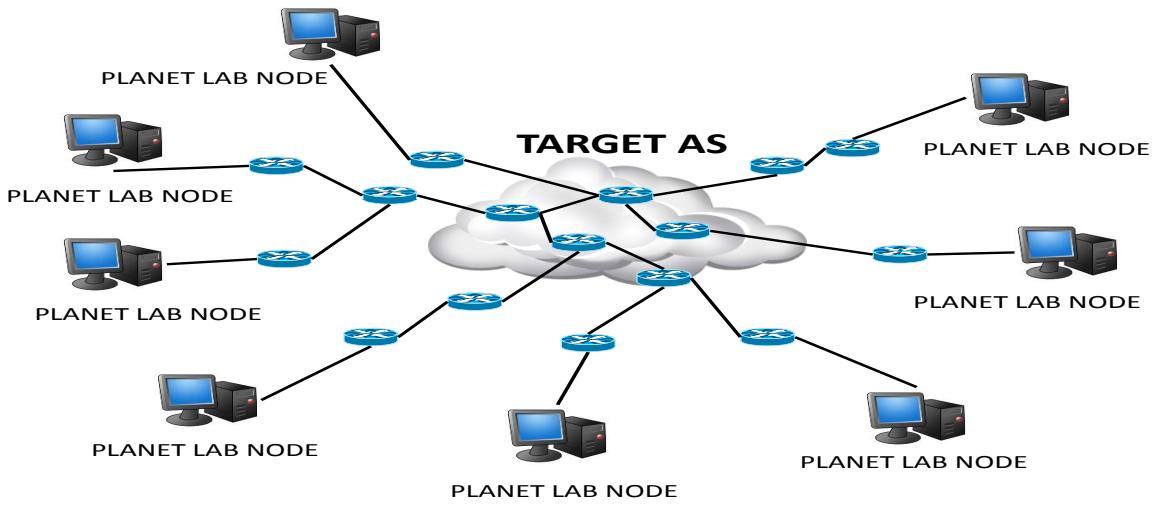

Figure 3.4: Traceroute from various planet lab nodes to various ip addresses residing inside target AS

3.2 Intra-AS router connectivity estimation

We next generate connectivity maps of routers within each AS leveraging the Rocketfuel tool [17], which relies on traceroute probes from looking glass servers across the globe, directed towards ASes that we wanted to study. However, as mentioned earlier, we used ∼500 PlanetLab [4] to execute Rocketfuel [17] for inferring the network connectivity inside the ASes. From these ∼500 PlanetLab machines distributed across the globe, we launched multiple traceroute probes to multiple IP addresses of target AS/System. Figure 3.4 depicts how traceroute from multiple locations is used to map router-level topology of target AS.

We additionally generate network maps showing the connectivity of routers leading to several important networked installations within the country, such as important educational institutions, banks and financial organizations, military and defense websites, research organizations and government networks.

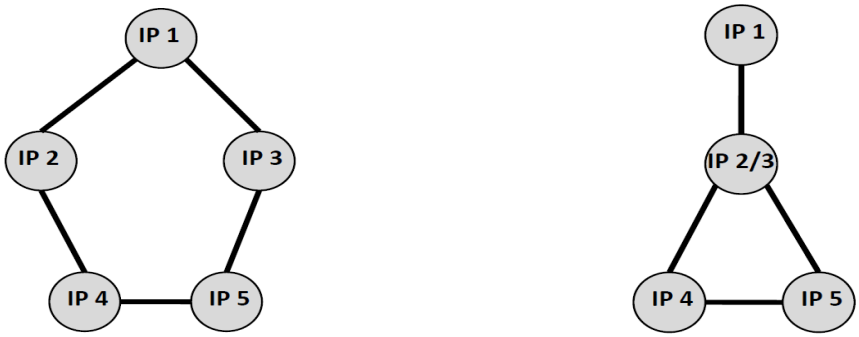

(a) Topology before Alias Resolution

(b) Topology after Alias Resolution

Figure 3.5: Alias Resolution

ALIAS RESOLUTION : The inbound and outbound interfaces of a router often show up as two different routers in traceroute. Thus, accurate router level topology from traceroute data involves *alias* resolution. There are various alias resolution tools that can associate the different router interfaces to the same one [11, 12] . In our research, we leverage Midar [12], which uses the classical approach involving IP ID field values to determine if two IP interfaces belong to the same router or not. Figure 3.5 demonstrates this problem and the corresponding solution. After having resolved the aliases on the traceroute results, we applied peer centrality [2] to determine routers that appear in several paths simultaneously.

Chapter 4

Evaluation and Results

This section is organized as follows :-

1.) Presenting the results of mapping Indian ASes and their connectivity followed by some statistics with respect to each AS, such as number of hosts that are potentially affected if an AS in compromised or unavailable.

2.)Results representing the intra-AS router level connectivity information of top three ASes of India.

3.)Results of mapping routers leading to important institutions in the country such as banks, educational institutions, hospitals, defence websites etc.

4.1 Autonomous Systems(ASes) that potentially transport large fraction of Indias Internet traffic

Figure 4.1 shows the top autonomous systems in India that intercepts major fraction of India's entire Internet traffic. These ASes are identified by applying peer centrality [2] on the AS connectivity graph generated using Gao *et al* MULTIPATHS algorithm [15].

Figure 4.2 shows the graph between top ASes in India and fraction of traffic transiting those ASes. The graph can be interpreted as :- 46.87% of traffic transit through AS9498(rank-1), 55.85% of traffic transit through AS9498(rank-1) and AS17439(rank-2) combined, 76.62% of traffic transit through AS9498(rank-1), AS17439(rank-2) and AS4755(rank-3) combined and so on. Thus, it is evident from the graph that 99.61% of inferred traffic paths crosses these 12 ASes that are potentially both customer as well as transit ASes.

Fig 4.3 presents the potential number of hosts behing each AS that are affected when each of those ASes is compromised or unavailable. The potential number of hosts that are affected by each AS are the sum total of the potential number of hosts owned by that particular AS and potential number of hosts owned by each of its singly owned customers(and their subsequent singly homed customers).

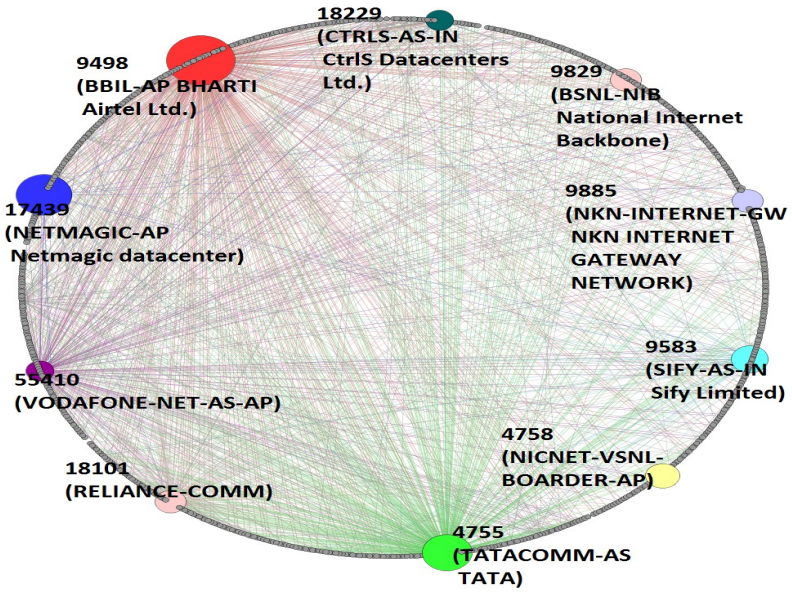

Figure 4.1: AS-connectivity graph of India with top/major ASes

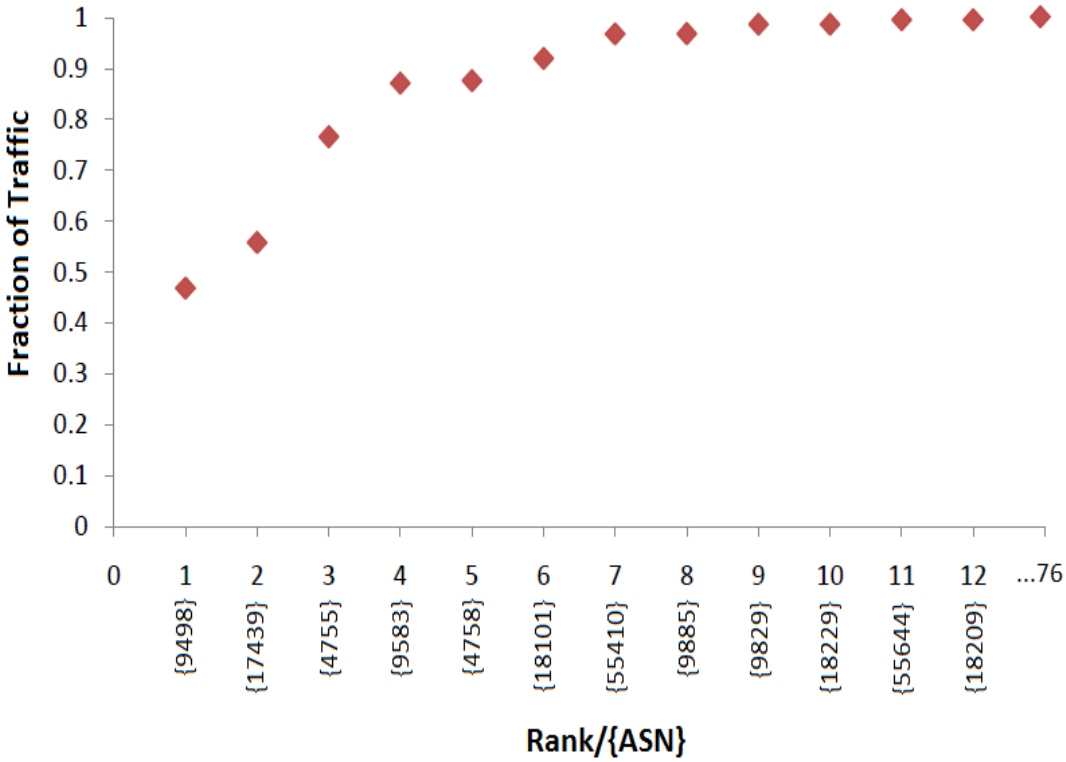

Figure 4.2: Fraction of paths through major ASes

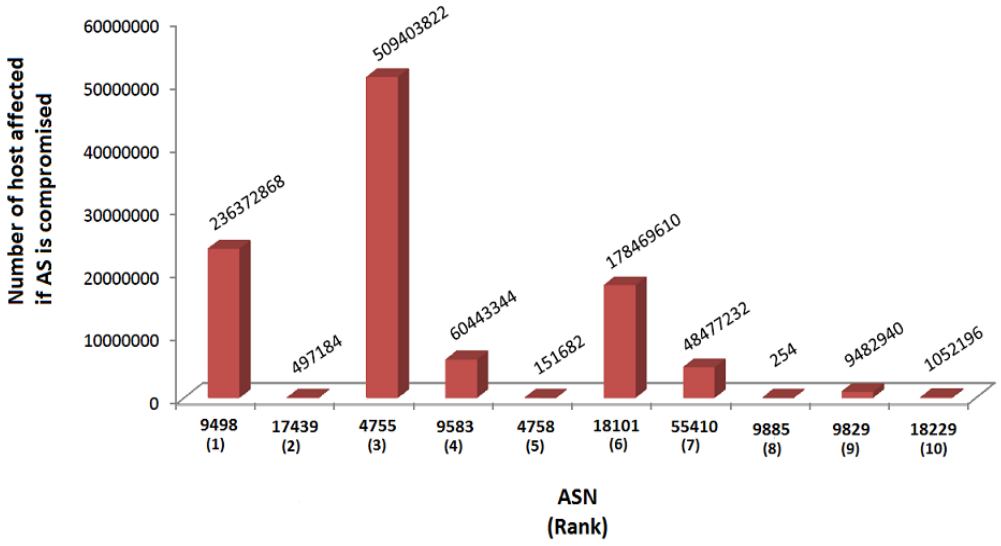

Figure 4.3: Number of hosts behind each ASes

However you can see that critical rank-2 AS comparatively affects less number of hosts than critical rank-3 and 4 in case of a mishappening which is bit strange, this is because the centrality is derived from number of paths that traverse the ASes. However, the number of hosts is derived from the AS prefix information. Critical rank-2 AS i.e AS17439(netmagic datacenter) peers with large number of ASes but owns less number of hosts(has less prefixes) and has few singly homed customer ASes.

4.2 Interdomain results(Router-connectivity results)

In the previous subsection we presented results of our efforts to identify the ASes which transport large fraction of India's Internet traffic. The top three ASes, ranked on number of paths to various ASes, happens to be Airtel(AS9498), Netmagic(AS17439) and Tata Communications(AS4755). We now present results of mapping the routers inside each of the top three ASes. Figure 4.4, 4.5 and 4.6 present the result as a cumulative distribution of the amount of traceroute flows that the routers in each AS carry. To identify these routers, we emulated the Rocketfuel [17] strategy that involves traceroute to various IPs in each network. The paths are thereafter input to alias resolution tools to uniquely identify routers in the network.

The x-axis of the graph presents the critical routers in decreasing order of number of traceroute paths that traverse it. For e.g. maximum number of traceroute flows pass through router number 1, followed by router number 2 and so on. The actual router IPs/DNS names are being suppressed for the sake of privacy and secrecy. We thus refer to them with their ranks. From the perspective of vulnerability to external attacks, these router are the critical ones as they lie

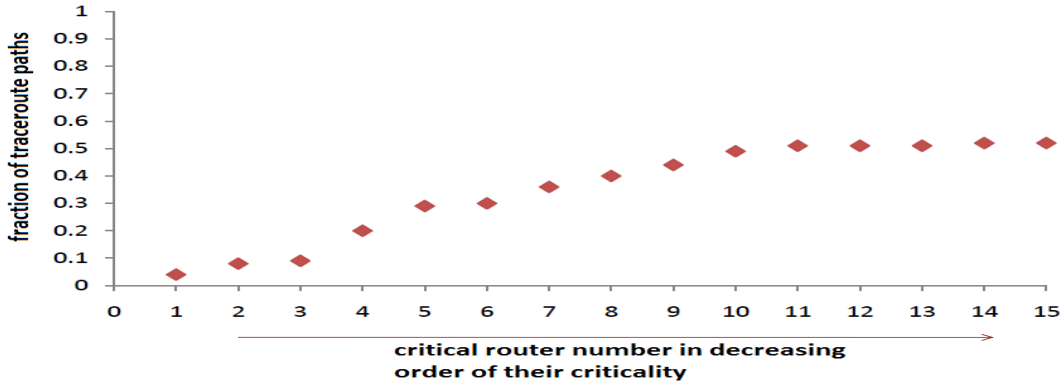

Figure 4.4: Graph between number of routers vs the number of traceroute flows that pass through them (for Airtel Network)

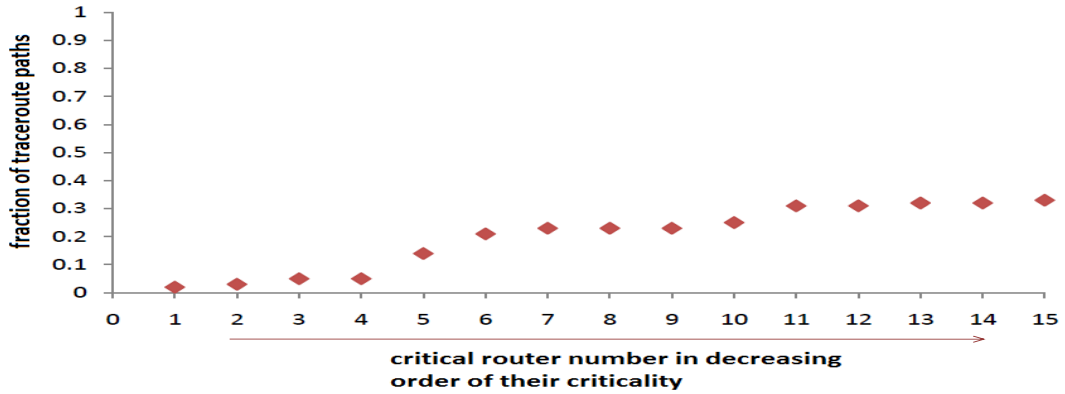

Figure 4.5: Graph between number of routers vs the number of traceroute flows that pass through them (for Netmagic Network)

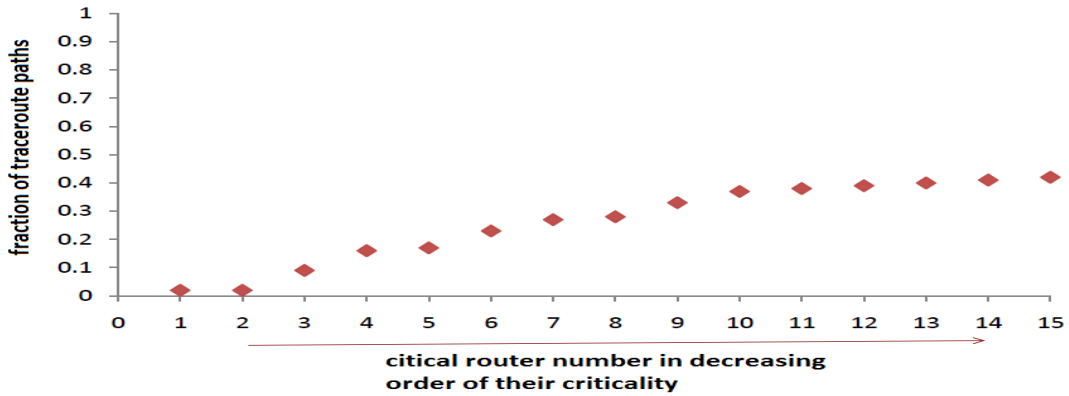

Figure 4.6: Graph between number of routers vs the number of traceroute flows that pass through them (for Tata Comm. Network)

at the intersection of several paths (potentially carrying lot of traffic of the network).

The y-axis of the graph represents the fraction of traceroute paths that transit through these routers. For e.g., in Airtel network, 4.96% of all traceroute paths transit through router 1. 8.05% of all traceroute paths transit through router 1 and 2 together. The 15 routers together transport 52.68% of the traceroute paths. Similarly, in case of Netmagic and Tata Communications these figures are 33.74% and 40.4% respectively. However, it must be kept in mind that these routers are the ones that potentially intercept large fraction of network path. These figures do not represent the actual fraction of the traffic that transport these routers.

4.3 Identifying routers that transport large fraction of traffic to various important sites

Similar to the aforementioned efforts, we tried to identify routers that transport large fraction of traffic to various important destination within India e.g. important government sites, defense websites, banks etc. The process for identifying such routers was very similar to the one used for identifying the important routers in each AS. We emulated the Rocketfuel [17] technique that involved launching traceroute from various PlanetLab [4] sites to the chosen destinations. Thereafter, alias resolution was applied to identify unique routers in each of the paths that were output by traceroute. Finally, centrality [2] approach was applied to each path to identify the critical routers that transported large number of traceroute flows. The results of these findings are presented in tables 4.1, 4.2, 4.3, 4.4, 4.5, 4.7.

In these tables, the first column represents the AS that hosts each of the target site. The second column represents the number of critical routers (and their corresponding ASes) that transport large fraction of the traceroute flows to the destinations on day1 dated 20/4/15. The third column presents the percentage of traceroute flows that pass through each of those routers respectively. The fourth column represent the number of critical routers(and their corresponding ASes) that transport large fraction of traceroute flows to the destination on day2 dated 15/6/15. Finally, fifth column presents the percentage of traceroute flows that pass through each of these routers respectively on day2.

In some cases, such as popular E-commerce sites, we observed that most of these popular sites do not have their servers inside India but hosted outside India in popular cloud storage services thats why the number of critical routers and traffic transiting through these critical routers are 0 as we are completely focused on finding the critical routers within Indian boundaries. A similar pattern was also seen for websites of hospitals.

These 2 sets of experiments to identify critical routers that transport large fraction of traffic to various important sites were conducted in an internal gap of approximately one and half months, but the results were pretty much the same, IP addresses of critical routers after 2 months have not been changed. Thus, it is evident from the results that critical routers information are static and does not change over a period of 2 to 3 months.

ASN which host target site	**# of critical routers on day 1**	**% of traceroute flows through critical routers on day 1**	**# of critical routers on day 2**	**% of traceroute flows through critical routers on day 2**
AS55479	2{18101,4755}	54.76,49.86	2{18101,4755}	52.68,53.49
AS55824	1{55824}	100	1{55824}	100
AS31815	0	0	0	0
AS2697	2{2697,4755}	97.04,73.96	3{2697,4755,9498}	97.26,72.13,64.48
AS55824	1{4755}	100	1{4755}	100
AS36351	0	0	0	0
AS55824	1{55824}	97.53	1{55824}	100
AS55847	1{4755}	100	1{4755}	100
AS9829	2{9829,9829}	100,60.58	1{9829}	100
AS46606	0	0	0	0

Table 4.1: Number of *critical* routers and their ASes that transport large number of paths leading to important educational institutions

ASN which host target site	**# of critical routers on day 1**	**% of traceroute flows through critical routers on day 1**	**# of critical routers on day 2**	**% of traceroute flows through critical routers on day 2**
AS9583	1{9583}	100	1{9583}	100
AS4755	2{4755,4755}	100,97.85	2{4755,4755}	100,97.43
AS17436	1{9498}	59.4	1{9498}	55.31
AS4755	1{4755}	99.37	1{4755}	98.84
AS4755	1{9498}	95.68	1{9498}	69.92
AS133708	2{133708,133708}	51.34,50.74	2{133708,133708}	49.86,49.6
AS20940	1{4755}	100	1{4755}	100
AS17439	2{17439,4755}	100,98.13	3{4755,17439,9498}	97.70,97.70,95.97
AS9238	3{9238,4755,9238}	100,95.29,94.70	3{9238,9238,4755}	100,95.08,95.08
AS9583	2{9583,9583}	100,94.11	AS9583 2{9583,9583}	100,91.80

Table 4.2: Number of *critical* routers and their ASes that transport large number of paths leading to important financial institutions(*e.g.* banks)

ASN which host target site	**# of critical routers on day 1**	**% of traceroute flows through critical routers on day 1**	**# of critical routers on day 2**	**% of traceroute flows through critical routers on day 2**
AS4758	1{9498}	59.57	1{9498}	60.41
AS4755	2{4755,4755}	100,97.85	3{4755,4755}	100,97.43
AS4758	1{9498}	100	1{9498}	97.82
AS4758	1{9498}	95.65	1{9498}	100
AS4758	1{9498}	63.82	1{9498}	65.3
AS17813	1{4755}	99.44	2{4755,17813}	100,62.13
AS4758	1{4755}	97.67	2{4755,4758}	100,66.66
AS4758	1{9498}	93.47	1{9498}	100
AS4758	1{9498}	100	1{9498}	97.77
AS4758	1{4755}	97.61	AS4758 2{4755,4758}	100,65.11

Table 4.3: Number of *critical* routers and their ASes that transport large number of paths leading to important Indian Government institutions(*e.g.* Revenue Services Department, Prime Ministers Office, Railways and Transport Department)

ASN which host target site	# of critical routers on day 1	% of traceroute flows through critical routers on day 1	# of critical routers on day 2	% of traceroute flows through critical routers on day 2
AS9752	1{17439}	28.2	1{17439}	30.79
AS16509	1{9498}	91.3	1{9498}	90.62
AS17439	0	0.0	0	0.0
AS38895	0	0.0	0	0.0
AS38895	0	0.0	0	0.0
AS11643	0	0.0	0	0.0
AS16625	0	0.0	0	0.0
AS38895	2{18229,4755}	98.51,99.25	2{18229,4755}	99.47,98.43
AS18229	0	0.0	0	0.0
AS40034	0	0.0	AS20940 0	0.0

Table 4.4: Number of *critical* routers that transport large number of paths leading to popular E-commerce sites (used by Indians) hosted in India.

ASN which host target site	# of critical routers on day 1	% of traceroute flows through critical routers on day 1	# of critical routers on day 2	% of traceroute flows through critical routers on day 2
AS4758	1{4755}	100	1{4755}	100
AS4758	1{9498}	63.82	1{9498}	64.58
AS4758	1{4755}	100	1{4755}	100
AS4758	1{4755}	100	1{4755}	100
AS4758	1{4755}	100	1{4755}	98.85
AS4758	1{4755}	100	1{4755}	100
AS4758	1{4755}	100	1{4755}	100
AS4758	1{9498}	100	1{9498}	100
AS4758	2{4755,4758}	100,78.94	2{4755,4758}	100,70
AS4758	1{4755}	97.27	AS4758 1{4755}	97.72

Table 4.5: Number of *critical* routers that transport large number of paths to important defense websites

ASN which host target site	# of critical routers on day 1	% of traceroute flows through critical routers on day 1	# of critical routers on day 2	% of traceroute flows through critical routers on day 2
AS4758	1{9498}	95.65	1{9498}	97.91
AS4758	1{4755}	100	1{4755}	100
AS4758	1{9498}	97.77	1{9498}	97.87
AS4755	2{4755,4755}	99.43,97.74	3{4755,4755,4755}	100,93.95,93.40
AS18101	1{18101}	94.67	2{18101,18101}	96.17,95.90
AS18101	1{18101}	94.67	2{18101,18101}	96.17,95.90
AS132215	1{18101}	73.46	1{18101}	71.51
AS4758	1{9498}	63.82	1{9498}	64.70
AS4758	2{4755,4755}	99.38,96.29	2{4755,4755}	99.42,98.27
AS133256	3{133256,133256,9583}	99.41,93.52,91.76	AS133256 3{133256,9583,133256}	100,92.89,90.71

Table 4.6: Number of *critical* routers that transport large number of paths to important power generation and distribution sites.

ASN which host target site	# of critical routers on day 1	% of traceroute flows through critical routers on day 1	# of critical routers on day 2	% of traceroute flows through critical routers on day 2
AS30083	0	0.0	0	0.0
AS8560	0	0.0	0	0.0
AS32244	0	0.0	0	0.0
AS4758	1{9498}	56.00	1{9498}	58.18
AS40034	0	0.0	0	0.0
AS33480	1{55410}	78.75	1{55410}	70.19
AS55824	1{55824}	99.39	1{55824}	98.94
AS33480	1{55410}	74.39	1{55410}	74.25
AS29873	0	0.0	0	0.0
AS20773	0	0.0	AS20773 0	0.0

Table 4.7: Number of *critical* routers that transport large number of paths to websites of well known hospitals

The another important aspect of these results is that none of the critical routers of airtel(AS9498) and netmagic(AS17439) from section 4.2 are the ones that appear as critical routers leading to specific destinations(educational institutions websites, financial institutions, Indian government institutions, e-commerce websites, defense websites,hospital websites, power generation and distribution websites) whereas there are few critical routers of tata(AS4755) that appear as critical routers leading to financial institutions(banks) and power generation and distribution websites.

Chapter 5

Discussion

The main motivation for our research was to obtain an insight into the topology of networks in India. To that end we implemented Gao *et al* [15] AS path mapping algorithm to determine the ASes inside India and their connectivity information. Their algorithm performs a depth first search to map the path from a prefix to all ASes. As described previously, we used all the prefixes listed in the Alexa database [1], corresponding to websites that are accessed by users in India. The list of ASes are obtained from the routing tables obtained from the Routeviews database. Gao *et al* [15] algorithm essentially performs a depth first search rooted at each of the input prefixes. All the experiments were performed on machine equipped with Intel Core i3 processor and 1GB RAM.

The first phase of the experiments provides us with information related the connectivity of ASes inside India. Our network centrality computation was applied to this graph. Thereafter, we obtained the ASes that cumulatively transport a very large fraction of India's Internet traffic. It is very evident from the results that a very small fraction of ASes (12 ASes) together transports a large fraction (99.61%) of India's Internet traffic.

The second phase of the experiments involves peering inside these ASes to determine key network routers that transport large fraction of traffic for the AS. This required us to map inside the network. This was achieved by emulating the Rocketfuel [17] strategy, as discussed earlier. In the second phase most of the results were gathered using traceroute probes with UDP, ICMP and TCP. However, we observed that majority routers in India blocked ICMP and TCP probes. Our results thus used the information available mostly from regular UDP based probes. Even though UDP traceroute probes were not filtered by most routers, there were still several routers which didn't decrement the TTL hop count. We thus failed to determine their IP addresses.

We mapped the internal networks of top three ASes. Traceroutes from various planet lab nodes provides us with information related to connectivity of routers inside the targeted AS. Then this graph is fed into centrality computation algorithm which gives us the routers which transported the major fraction of traffic for that ASes(hence these routers are the critical ones).

5.1 Security Considerations

A part of our research motivation was to answer questions like – *Does there exists a small set of ASes and routers through that transports large fractions of India's Internet traffic?*. The answer to such questions opens up avenues for more challenging security questions like – *Is it easy for an attacker to take bring down large fraction of network connectivity by merely attacking a small number of networks and devices?* . Moreover, it may aid an adversary to not only attack large fraction of network traffic but might also make it easy for the attacker to monitor (and perhaps censor) large fraction of the network traffic. Our research is the first step towards answering such questions.

From our results it is evident that there exist a very small fraction of ASes (12 ASes out of 800 ASES) together transports a large fraction (99.61) of India's Internet trac. We also found that there are some critical routers on the path of important website destinations through which major fraction of traffic flows through. Hence an adversary can easily snoop on these one of the critical routers or more and listen to major fraction of traffic without snooping on many devices simultaneously. Therefore, extra theft protection must be given to these routers and ASes as they are the critical ones.

Chapter 6

Conclusion

In this paper, we have tried to measure/discover the network topology of India's entire Internet space. India which is the second largest nation by population has approximately 250 million Internet users. With such large number of Internet users in India, there are only 76 ASes which connects India's internet to the world outside which is relatively very less compared to other developed nations like USA. We explained that the Internet topology discovery is driven by important questions such as - *is it possible for an attacker to take down large fraction of network connectivity by merely attacking small number of networks and devices?*

So, at first we took leverage from Gao *et al* [15] algorithm to map all the ASes in India and then with peer centrality we are able to locate the critical/vulnerable ASes of India. Our result shows that 99.61% of Internet traffic flows through top 12 ASes of India.

Second we emulated rocketfuel tool [17] to map all the routers inside critical ASes and routers leading to several important networked installations within the country.

There are many interesting applications that can potentially benefit from the network topology discovery and measurement such as *network diagnosis, performance optimization and reliability enhancement in multihoming, content distribution* and *peer-to-peer* applications.

Summary of Contributions

- Developed the AS connectivity graph of India and figured out some crucial statistics based on that graph.
- Mapped the router connectivity graph of major Internet players in India and based on that graph figured out some vital information about them.
- Developed the network maps showing the connectivity of routers leading to several important networked installations within the country and identification of critical routers leading to those networked installations

Future Work : For future work, we would collaborate with ISPs to determine what level of information is available and use our results as to study the actual engineering efforts that an attacker needs to employ to launch various kinds of attacks (Denial of Service, Eavesdropping etc.) on a large fraction of the network traffic.

Bibliography

[1] Alexa-actionable analytics for the web.

[2] Centrality-https://en.wikipedia.org/wiki/centrality.

[3] Hurricane electric bgp toolkit-www.bgp.he.net.

[4] Planetlab http://www.planet-lab.org/.

[5] Route views project, university of oregon-routeviews.org.

[6] Branigan, S., Burch, H., Cheswick, B., and Wojcik, F. What can you do with traceroute? *Internet Computing, IEEE 5*, 5 (2001), 96.

[7] Donnet, B., and Friedman, T. Internet topology discovery: a survey. *Communications Surveys & Tutorials, IEEE 9*, 4 (2007), 56–69.

[8] Faloutsos, M., Faloutsos, P., and Faloutsos, C. On power-law relationships of the internet topology. In *ACM SIGCOMM computer communication review* (1999), vol. 29, ACM, pp. 251–262.

[9] Gao, L. On inferring autonomous system relationships in the internet. *IEEE/ACM Transactions on Networking (ToN) 9*, 6 (2001), 733–745.

[10] Govindan, R., and Tangmunarunkit, H. Heuristics for internet map discovery. In *INFOCOM 2000. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE* (2000), vol. 3, IEEE, pp. 1371–1380.

[11] Gunes, M. H., and Sarac, K. Analytical ip alias resolution. In *Communications, 2006. ICC'06. IEEE International Conference on* (2006), vol. 1, IEEE, pp. 459–464.

[12] Keys, K., Hyun, Y., Luckie, M., and Claffy, K. Internet-scale ipv4 alias resolution with midar. *IEEE/ACM Transactions on Networking (TON) 21*, 2 (2013), 383–399.

[13] Koradia, Z., Mannava, G., Raman, A., Aggarwal, G., Ribeiro, V., Seth, A., Ardon, S., Mahanti, A., and Triukose, S. First impressions on the state of cellular data connectivity in india. In *Proceedings of the 4th Annual Symposium on Computing for Development* (2013), ACM, p. 3.

[14] McRobb, D., Claffy, K., and Monk, T. Skitter: Caidas macroscopic internet topology discovery and tracking tool, 1999.

[15] Qiu, J., and Gao, L. As path inference by exploiting known as paths. In *Proceedings of IEEE GLOBECOM* (2005).

[16] Shavitt, Y., and Shir, E. Dimes: Let the internet measure itself. *ACM SIGCOMM Computer Communication Review 35*, 5 (2005), 71–74.

[17] Spring, N., Mahajan, R., and Wetherall, D. Measuring isp topologies with rocketfuel. In *ACM SIGCOMM Computer Communication Review* (2002), vol. 32, ACM, pp. 133–145.

[18] Spring, N. T., Wetherall, D., and Anderson, T. E. Scriptroute: A public internet measurement facility. In *USENIX Symposium on Internet Technologies and Systems* (2003).