Can We Study the Topology of the Internet from the Vantage Points of Large and Small Content Providers?

by

Aleisa A. Drivere

Submitted in Partial Fulfillment of the Requirements

for the Degree of

Master of Computing and Information Systems

in the

Computer Science and Information Systems

Program

YOUNGSTOWN STATE UNIVERSITY

December 2011

Can We Study the Topology of the Internet from the Vantage Points of Large and Small Content Providers?

Aleisa A. Drivere

I hereby release this thesis to the public. I understand that this thesis will be made available from the OhioLINK ETD Center and the Maag Library Circulation Desk for public access. I also authorize the University or other individuals to make copies of this thesis as needed for scholarly research.

Signature:

Aleisa A. Drivere, Student Date

Approvals:

Dr. Graciela Perera, Thesis Advisor Date

Dr. Bonita Sharif, Committee Member Date

Dr. John R. Sullins, Committee Member Date

Dr. Peter J. Kasvinsky, Dean of School of Graduate Studies and Research Date

Abstract

The purpose of this study is to gather and analyze Internet data at the autonomous system level from the vantage points of large and small content providers, to see if any differences can be detected. Data was gathered using the traceroute program, and individual autonomous systems (ASes) were identified using BGP Looking Glass, and Whois databases provided by ICANN regional registries. The data was analyzed using the following metrics: Average AS path length, number of networks to which a content provider was connected, average number of major networks per path, average number of Internet exchange points (IXPs) per path, and geographical characteristics which categorize the ASes in each path based on IANA regional boundaries. The results show some interesting trends in routing and business relationships that will surely have an impact on the future Internet.

Acknowledgements

First and most importantly, I would like to thank my family, especially my husband Pete for his love, support, and for believing in me. Without him, none of this would have been possible

I would like to take the opportunity to thank my advisor Dr. Graciela Perera for her invaluable help and advise in matters both technical and professional, and for encouraging and believing in me. I would also like to thank my committee members Dr. John R. Sullins and Dr. Bonita Sharif for their editorship and suggestions.

In addition, I would like to thank Antoine Ogbonna, Sandeep Gundla, and John Mela for help with my research but more importantly the friendship and encouragement that made a difference. I would also like to thank Kyle Marlin for help with the data collection, and all of the faculty and staff in the Computer Science and Information Systems Department.

Table of Contents

Chapter 1

Introduction

The topology of the Internet is a complex area of study that only increases in complexity with time, along with the Internet itself. From four hosts in 1969 to approximately 100 in 1975, by 2010 there were over 700 million Internet-connected hosts, and that number promises to increase. With the Internet growing at steady rate, Internet topology research continues to be an area of active study. While there has been a great deal of research on the connectivity patterns of the Internet at the autonomous system level, this study analyzes Internet topological data from the *specific* vantage points of large and small Internet content providers, in an attempt to discover if any differences can be seen. What this study contributes is really an identification of *trends* in Internet routing from the content provider point of view. The presence (or absence) of large Internet Service providers and Internet exchanges as well as the connectivity level of individual content providers, may indicate changes in the Internet landscape that are driven by both technological and business concerns. Characteristics used in the comparison analysis were: number of unique autonomous systems per path, number of unique networks to which a content provider is connected, number of Internet exchanges per path, and geographical characteristics based on IANA regional Internet registries. It is hoped that the vantage point-specific results of this study can contribute to the available research on Internet topology, with the goal of better understanding routing, content delivery, and business trends affecting the Internet today.

Figure 1. Internet Host Count, January 1994 - January 2010

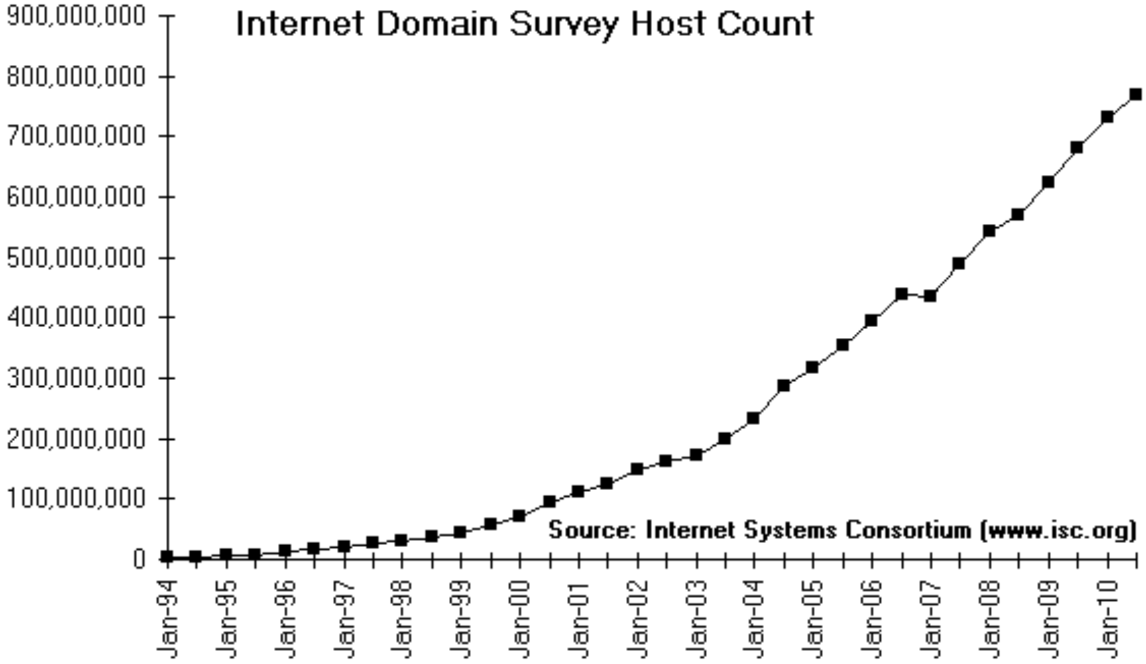

[1]

1.1 Historical Background

To attempt to understand the current structure of the Internet as well as the concepts in this study, it is best to begin at the beginning, when the Internet was a U.S. government research project consisting of only four connected systems. The story is well known and often repeated. In 1969, The Advanced Research Projects Agency (ARPA) allowed for the creation of a very small wide-area network (ARPANET) for research purposes. The first ARPANET nodes were located at the University of California, Los Angeles (UCLA), the Stanford Research Institute (SRI), the University of Utah, and the University of California, Santa Barbara. It was from this simple beginning that the Internet of today evolved, and each step in the evolution built upon the previous steps.

---

[1] (http://www.navigators.com/stats.html)

When the ARPANET was first designed, the decision was made to base its operation on packet switching, with data being transmitted through the network in discrete units (packets). By 1975, the ARPANET had about 100 nodes, and by the late 1970's, packet switching was being successfully used on wired networks as well as wireless networks connected by radio and satellite communication links. Another building block of the Internet was the Local Area Network (LAN), which consists of computers connected in a relatively small area such as a building or campus. A Wide Area Network (WAN) connects over a much larger geographical area. As the ARPANET grew in size and popularity, many LANs connected to it via gateway nodes. This created a *network of networks*; a term that even today is often used to describe the Internet. The ARPANET nodes formed a "backbone", while the LANs that connected to it enabled more and more end systems to remotely communicate with each other.

1.2 Internet Protocol

In 1983, the ARPANET implemented the Transmission Control Protocol/Internet Protocol (TCP/IP) in order to maintain a standard by which diverse and heterogeneous networks could easily communicate. Each node on an IP network has a unique *IP address*. Currently the most widely deployed IP addressing scheme is IP version 4 (IPv4) which uses a 32 bit address in "dotted-decimal notation", i.e. 150.134.8.71.[2] Blocks of IP addresses are assigned by the Internet Assigned Numbers Authority (IANA) to regional registries, which in turn assign blocks of addresses to individual netwoks. The IANA

---

[2] IP version 6 has been developed and deployed, although according to the Global IPv6 Deployment Progress Report (http://bgp.he.net/ipv6-progress-report.cgi), as of December 2, 2011 there were 34,870 ASes that are using only IPv4 as opposed to 4883 that are IPv6-capable.

regional registires are: American Registry of Internet Numbers (ARIN), European IP Network Coordination Center (RIPE NCC), Latin American and Caribbean Network Information Center (LACNIC), Asia Pacific Network Information Center (APNIC), and African Network Information Center (AfriNIC).

1.3 Autonomous Systems

In the 1980's the concept of the Autonomous System was introduced. An autonomous system (AS) is essentially a network or system of networks operating under a single administrative control, such as a company or a university. A single AS is often referred to as an *enterprise network*. Different ASes communicate with each other using the Border Gateway Protocol (BGP). The next section describes this key protocol, which allows connectivity between different ASes.

1.4 BGP

The Border Gateway Protocol version 4 (BGPv4 or simply BGP) is the *de facto* standard protocol for routing between different autonomous systems. BGP allows ASes to exchange information about their connected networks so that networks in individual ASes can reach every part of the Internet. BGP identifies each autonomous system by its globally unique autonomous system numbers (ASN). IANA regional registries assign ASNs. When a network is advertised through BGP, its IP prefix is included along with certain BGP *attributes*. Two of the most important attributes are AS-PATH and NEXT-HOP. The AS-PATH attribute contains the ASNs of the ASs through which a prefix has passed. The NEXT-HOP attribute is primarily of interest to the networks internal to an autonomous system, as it contains the IP address of the router interface that begins the

AS-PATH. This attribute provides a link between inter-AS and intra-AS routing. Given that BGP enables connectivity between different ASes, it allowed the Internet to evolve to a hierarchical structure explained in the next section.[3]

1.5 Hierarchical Structure of the Internet

In the late 1980's, the National Science Foundation began constructing the NSFNET in order to provide reinforcement and additional resources to the original ARPANET backbone infrastructure, and in 1990, the ARPANET was retired, leaving the NSFNET as the sole backbone network. The NSF encouraged the development of smaller regional networks in order to connect enterprise networks to the NSFNET. A hierarchical three-tiered structure began to emerge: NSFNET backbone, regional networks, and enterprise networks. By the early 1990's, many commercial Internet Service Providers (ISPs) were in business, and in 1995, the NSFNET backbone was retired. This left the Internet with no centralized backbone, as the commercial ISPs were all maintaining their own competing backbone networks.

To fill the Internet backbone void, the largest of these ISPs such as AT&T, Sprint, and Qwest, formed the new backbone of the Internet and became known as Tier-1 ISPs. Somewhat smaller ISPs and even smaller regional ISPs became known as Tier-2 and Tier-3, respectively. Although the classification of Tier-1, Tier-2, and Tier-3 is often used to describe ISPs and the hierarchical operation of the Internet, it is not always easy to determine the classification to which an ISP belongs. An ISP is typically identified as Tier-1 if it meets the following qualifications: 1) It is IP transit-free and 2) It does not pay any other network for peering. IP transit is where an ISP advertises its customers'

---

[3] (Kurose and Ross)

routes to other ISPs, and advertises a default route to its customers, in order to allow them to access the entire Internet. Peering is when two networks connect to each other so that traffic between the two peers and their customers can be exchanged. A transit-free network will peer with other networks, but does not provide any information about the routes of its customer networks. A Tier-2 network may peer with other networks, but still pays transit fees and/or peering settlements. Due to confidential business agreements, it is sometimes difficult to determine if a network pays for peering. Because of this, some large networks that are often identified as Tier-1 may actually be Tier-2. A Tier-2 ISPs will usually connect to one or more Tier-1 networks in order to reach every part of the Internet. Tier-3 networks, often called access networks, can only reach the Internet by paying for IP transit. It is through these Tier-3 ISPs that most end-users connect to the Internet. Large content providers such as popular search engines and e-commerce sites have also typically accessed the Internet via the Tier-1/Tier-2/Tier-3 hierarchical system, but there is strong evidence that this may be changing.[4]

1.6 Motivation and Related Work

A 2008 study done by Gill, Arlitt, Li, and Mahanti showed that large content providers like Microsoft, Yahoo, and Google are deploying their own wide area networks. One of the effects of such deployments is that traffic to and from these content providers tends to bypass major Tier-1 networks. As per Gill, et al., the reasons for this may include: 1) the vulnerability of smaller ISPs to de-peering and transit disputes involving Tier-1 and Tier-2 ISPs, hence an interruption of customer service. 2) Limitations and/or uncertainty about Internet content delivery due to issues such as IP multicast. With the emergence of

---

[4] (Gill, Arlitt and Li)

applications that involve the necessity of real-time content delivery, such as video-on-demand, large content providers may feel that it is more beneficial for their business model to avoid such uncertainty. 3) Plans by larger content providers to provide cloud computing services such as software-as-a-service to subscribing customers. Microsoft, Google and Yahoo are already offering such services, and the trend is likely to continue. Along with major content providers, content delivery networks such as Akamai and Limelight have been deploying their own WANs and offering services to smaller content providers who may not have the resources to deploy their own WANs. The straightforward and easily duplicated methodology as well as the interesting results of the Gill study provided an inspiration to do an independent verification study using the same methodology.[5] It also introduced the question of what other Internet topological characteristics might be discovered using similar methods. This study is attempting to discover if differences can be seen in the logical topology of the Internet from the vantage points of content providers like Microsoft, Yahoo and Google, compared to the vantage point of much smaller content providers. Metrics such as average AS path length, number of unique networks directly connected to each content provider, the average numbers of major networks and Internet exchanges per path, and the number of regional boundaries crossed in each path are used to attempt to quantify these differences.

Other related Internet topological studies include the ongoing projects by CAIDA[6] and the University of Oregon Route Views Project[7] to gather and analyze Internet connectivity at both the router level and AS level.

---

[5] (Drivere, Ogbonna and Gundla)

[6] (http://www.caida.org)

[7] (http://routeviews.org)

Chapter 2

Methodology

To obtain the topological data, the first step was to define a method for data collection. Next, choose the tools that would implement the method. Lastly, collect the data once the tools were correctly configured.

2.1 Method

The following method was developed to collect data for this study.

a. Identify a set of M "large" content providers and a corresponding set of M "small" content providers.
b. Choose N locations worldwide from which to issue traceroute queries. Traceroute is a popular Internet test measurement tool that is described in Section 4.1.1.
c. Issue M x N traceroute queries.
d. Parse the IP addresses in the traceroute output in order to identify the autonomous systems in the various paths.

2.2 Tools

To collect the data, the following tools were used:

2.2.1 Traceroute and DipZoom

Traceroute is an Internet test measurement tool that can run on any computer (host) that is connected to the Internet. The host issuing a traceroute query is known as the source. The source user issues a query by specifying a destination: either a hostname (i.e. www.ampreonrecorder.com), or an IP address. Traceroute then sends multiple, uniform size packets of data toward that destination. On the way to the destination, the packets pass through a series of routers. When a router receives one of these packets, it returns back to the source a short message giving its hostname and the IP address of the interface on which the probe was received. The distance between two directly connected routers is often called a hop. Traceroute is a customizable program, and there are different mechanisms used to issue queries. The original version sends UDP probes to high numbered destination ports. This is the method that is used by the versions of traceroute available through many UNIX-based operating systems, including Mac OSX, FreeBSD, and many distributions of Linux. Another method uses Internet Control Message Protocol (ICMP) echo request probes instead of UDP probes. This method is used in the version of traceroute available with various versions of Microsoft's Windows operating system. Yet another method uses TCP SYN probes to well-known ports such as the default port 80 used for web servers.[8] For the data collected in this study, Linux-based traceroute servers running IPv4 traceroute were selected, in an attempt to maintain uniformity in the traceroute method and output.

Deep Internet Performance Zoom (DipZoom) is a peer-to-peer Internet test measurement application developed at Case Western Reserve University. DipZoom relies on a network of volunteer nodes to provide a variety of popular Internet test measurement tools including ping, nslookup, dig and traceroute.[9] After selecting these tools, locations needed to be chosen from which to issue the traceroute queries. The next section discusses that process and rationale.

---

[8] (Luckie, Hyun and Huffaker)

[9] (http://dipzoom.case.edu)

2.2.2 Geographic locations for traceroute queries

Locations were selected based on two factors: a.) Worldwide Internet usage statistics obtained from Internet World Stats Usage and Population Statistics[10], and b.) Availability of DipZoom traceroute servers. According to Internet World Stats, as of March 2011 44% of Internet users were located in Asia, 10.3% were located in Latin America/Caribbean and 5.7% were located in Africa. However DipZoom traceroute server selection was much more limited in these locations than in North America and Europe. At the time of data collection, only China, Japan, Singapore and Taiwan were available in Asia, Argentina and Brazil were available in Latin America, and unfortunately none were available in Africa.

---

[10] (http://www.internetworldstats.com/stats.htm)

Europe/Middle East/Russia (RIPE):

1. Austria
2. Belgium
3. Switzerland
4. Czech Republic
5. Germany
6. Spain
7. Finland
8. France
9. Greece
10. Hungary
11. Italy
12. Netherlands
13. Poland
14. Sweden
15. UK
16. Israel
17. Jordan
18. Russia

North America (ARIN):

Canada:

19. Calgary
20. Ottawa
21. Vancouver
22. Waterloo

United States:

23. California
24. Texas
25. New York
26. Florida
27. Illinois
28. Pennsylvania
29. Ohio
30. Michigan
31. Georgia
32. North Carolina
33. New Jersey
34. Virginia
35. Colorado
36. Massachusetts
37. Indiana
38. Arizona
39. Tennessee
40. Missouri
41. Maryland
42. Wisconsin
43. Minnesota
44. South Carolina
45. Washington

Latin America/Caribbean (LACNIC):

46. Argentina
47. Brazil

Asia Pacific (APNIC):

48. Australia
49. China
50. Japan
51. Singapore
52. Taiwan

2.2.3 Ten Large Content Providers

These 10 were selected from the list of top 20 sites worldwide provided by the Web traffic analytics company Alexa.[11] Alexa has developed a toolbar that gathers information about Web traffic and site visits from toolbar users, and has devised a ranking system (TrafficRank) based on the results gathered. The 10 content providers were selected to reflect Internet usage patterns worldwide as well as the locations of the available DipZoom traceroute servers. Since roughly 10% of the DipZoom traceroute

---

[11] (http://www.alexa.com)

servers used were located in Asia (4 Asia locations out of 52 total), only one Asia-based content provider was included (Baidu.com). Selections were also made to eliminate redundancy in network ownership. Since some companies own multiple sites, only one site per company was included. Using this requirement, the following sites in Alexa's top 20 were not used: YouTube, Blogger, Google India, Google.hk and Google.de (Google), MSN (Microsoft), and Yahoo Japan (Yahoo). The profiles of the ten large content providers include popular search engines (Google, Yahoo and Baidu), social networking sites (Facebook, Twitter and LinkedIn), web applications and software-as-a-service (Live), e-commerce (Amazon), web publishing (Wordpress) and information resources (Wikipedia). See Table 1 for the list of large content providers.

2.2.4 Ten Small Content Providers

These 10 were selected at random in an attempt to reflect the same demographic as the 10 large content providers. The only requirement was that in order to qualify as "small", a content provider could not appear in the list of Alexa's top 500 sites. The profiles of the small content providers include an educational institution, Youngstown State University (YSU), websites for businesses and organizations (Ampreonrecorder, Popdetective, Endlessanalog, Novaroma), an online magazine (Mixonline), informational sites (Developphp, Italianfoodforever), student resources (6students), and an Indian/English news and entertainment portal (Liveidiots). See Table 1.

Table 1. Large and Small Content Providers

Large Content Providers	Small Content Providers
www.google.com	www.ysu.edu
www.facebook.com	www.ampreonrecorder.com
www.yahoo.com	www.popdetective.com
www.baidu.com	www.endlessanalog.com
www.live.com	www.developphp.com
www.wikipedia.org	www.italianfoodforever.com
www.twitter.com	www.novaroma.org
www.linkedin.com	www.mixonline.com
www.amazon.com	www.6students.com
www.wordpress.com	www.liveidiots.com

2.2.5 List of Major Networks

Due to the difficulty in identifying Tier-1 and Tier-2 networks as mentioned in Section 1.5, the networks used in this study that qualify as major networks are taken from the list of the top 10 networks supplied by the Cooperative Association for Internet Data Analysis (CAIDA). The networks are ranked according to the size of their customer cone (the number of ASes that connect to that network). Taken together, these networks connect to between 74% and 96% of the ASes worldwide. Table 2 shows the CAIDA top 10 networks as of January 2011 along with their AS numbers and customer cone size.[12]

---

[12] (http://as-rank.caida.org)

Table 2. CAIDA Top 10 networks, January 2011

Rank	AS Number	AS Name	Percentage of ASes (customer cone)	Number of ASes (customer cone)
1	3356	Level 3	96%	35,753
2	6939	Hurricane Electric	91%	33,621
3	3549	Global Crossing	90%	33,427
4	6461	Metromedia Fiber	82%	30,524
5	3257	Tinet SpA.	81%	29,989
6	1239	Sprint	77%	28,636
7	2914	NTT America	77%	28,501
8	174	Cogent/PSI	75%	27,722
9	1299	TeliaNet	74%	27,573
10	7018	AT&T	74%	27,375

2.2.6 BGP Looking Glass

BGP Looking Glass is an application that runs on a server that acts as a portal to the routers of the organization that maintains the server. Looking Glass gives read-only access to the routing table information of the participating routers, specifically the AS numbers in the AS-PATH for a network in the routing tables.[13] BGP Looking Glass applications made available by Telia Sonera International Carrier[14] and Portland Oregon

---

[13] (http://www.bgp4.as)

[14] (http://lg.telia.net)

based Internet Partners Inc.[15] were used to identify AS numbers based on the IPv4 addresses in the traceroute output.

2.2.7 Whois

The Whois databases made available by the various IANA regional Internet registries were used to identify networks by IP address, when AS number information was not available from BGP Looking Glass. These databases were also used to provide geographic information about the ASes.

---

[15] (http://whois.ipinc.net/cgi-bin/lg.pl)

2.3 Data Collection

2.3.1 Collection Procedure

Traceroute queries were issued from 52 geographically distributed traceroute servers, to content providers' servers. The traceroute servers were made available through the DipZoom application. The locations of the traceroute servers are listed in Section 2.2.2. A single traceroute query was issued to each content provider from each traceroute server using the content provider hostname. Hostnames were used to allow DNS resolution local to the source. This was done to produce as accurate a snapshot as possible of real Internet traffic, given the limitations of traceroute.[16]Data was collected from May through September 2011.

2.3.2 Internet Exchange Points

BGP Looking Glass was not always able to return the AS numbers of some of the IP addresses in the traceroute output.[17] Upon further investigation it was found that many of these "unidentified" IP addresses belonged to Internet Exchange points. An Internet Exchange Point (IX or IXP) also sometimes called a Network Access Point (NAP) allows ISPs to exchange traffic between their networks by means of mutual peering agreements. The networks connected to an IX can reduce their reliance on 3rd party transit networks, with the primary goals of reducing cost and latency, and increasing bandwidth and fault tolerance.[18] This type of agreement between an IX and its connected networks is similar to the peering agreements between large so-called Tier-1 ISPs. As only routers in

---

[16] (Crovella and Krishnamurthy)

[17] When AS numbers were not available through Looking Glass, the organization identifiers (OrgID) given in the appropriate regional Internet registries were used.

[18] (http://www.bgp4.as)

networks connected to an IX would know that IX's AS number, this explains why the BGP routing tables of the Looking Glass servers used were not often able to provide AS numbers. When the presence of these IXs was observed, this provided an additional metric with which to analyze the data: average number of IXs per path, per content provider. A list of Internet Exchange Points found at http://www.bgp4.as/internet-exchanges was used to identify IX points in the traceroute output. Table 3 shows the Internet Exchange points that were identified in the collected data. Some of these IXs such as Hong Kong IX, Amsterdam IX, and London IX are government-operated infrastructures located in one geographical location, but some IXs are operated by commercial entities and the IP data and Whois information on them indicates a more widely deployed infrastructure.[19] Because it is expected that the AS numbers of IXs will be present in the routing tables of routers in connected networks, this list may not include all of the IXs in the collected data. Nevertheless, it is argued that analyzing the data with these IXs will still provide a good indication of the IX presence in each content provider's topology.

---

[19] Equinix Inc., Peer 1 Network Inc., Telehouse International Corp. and Pacific Wave are some examples.

Table 3. Internet Exchange Points

Name	Location	Organization ID
Equinix Inc.	United States	EQUINIX
TerreNAP Data Centers	United States	TERRENAP
Amsterdam IX	Netherlands	AMS-IX
Service for French IX	France	SFINX
Swiss IX	Switzerland	SWISSIX
Seattle IX Inc.	United States	SEATT-11
Peer 1 Network Inc.	Canada	PER1
Hong Kong IX	Hong Kong	HKIX
Netnod Swedish IX	Sweden	NETNOD
Telehouse Int'l Corp.	United States	TICA
CoreSite	United States	COWIL
Pacific Wave	United States	PACIFIC-24
Milan IX	Italy	MIX-NOC
London IX	United Kingdom	LINX
MAGPI Research	United States	MAGPI

2.3.3 Challenges and Limitations affecting the Data Collection Process

Some factors affecting the data collection should be noted. During the data collection process, it was observed that the majority of the traceroute servers available through DipZoom were affiliated with Planet Lab. As such, some of the challenges presented by

Planet Lab also affected the data collected for this study, primarily the availability of particular nodes[20]. When a particular node that had been used previously was discovered to be unavailable, a replacement node was selected that had the same IP prefix as the original node. If that was not possible, a replacement was selected that was located in the same city. Another factor to mention is that the majority of the traceroute servers were located at colleges and universities, which introduces some bias into the data. As the goal is to show a comparison in Internet topological characteristics, it is argued that this bias does not present a problem since it is consistent for all data collected. Also, when the traceroute output was examined, it was observed that some of the content provider networks did not return router information. This often happens for example, when network administrators disable ICMP *echo reply,* or restrict incoming UDP traffic. In traces where this happened, the information from the last router to return a response was examined, to determine if it was a border router to the content provider's network. Only when this determination could be made was the content provider's network assumed to be the next autonomous system. Lastly, it is recognized that traceroute data only provides information about the *forward* IP path. Therefore, no assumptions can be made about the *return* routing path from destination to source.

2.4 Data Analysis

After data was collected, it was analyzed using five metrics: Average AS path length, number of directly connected networks, average number of major networks per path, average number of Internet exchanges per path, and geographical boundaries crossed. There were some challenges in the data analysis, which are described in Section 2.4.2.

---

[20] (Peterson, Pai and Spring)

2.4.1 Metrics

The following metrics were used to analyze the data. These metrics were chosen in an attempt to reveal differences in Internet topology from content provider vantage points, while avoiding as much as possible some of the issues related to traceroute-based topological data.[21]

1. **Average AS-path length (APL):** This is the average number of unique autonomous systems in each path from source to destination. This metric gives an indication of how many different networks may route Internet traffic to the particular content provider.
2. **Number of Connected Networks (NCN):** The number of unique networks to which a content provider is directly connected. This metric is used as an indication of the potential *extent* of a content provider's own network. The higher the number, the more connected a content provider's network is worldwide.
3. **Average Major Networks per path (AMN):** This metric is used to characterize the *nature* of the content provider's Internet connectivity. The nature of the connectivity is determined by how closely it compares to traditional Tier-1/Tier-2/Tier-3 Internet connectivity patterns.
4. **Average number of Internet Exchanges per path (AIX):** This is an indication of how likely it is that traffic from worldwide sources to the content provider network will pass through an Internet exchange point.
5. **Number of Geographical Boundaries crossed in each path (GEOB):** The various IANA regional Internet registries used in the data collection process categorize the geographical boundaries. This metric is intended to identify some

---

[21] (Amini, Shaikh and Schulzrinne)

of the geographical characteristics of the traffic to a particular content provider. When a traceroute query is issued to a content provider, the expected result would be that the geographic locations of the ASes in the path would reflect the geographic locations of the source and the destination. The results of analysis with this metric will note exceptions to the expected result. While it is recognized that the geographical landscape of Internet traffic is extremely complex, involving many international carriers and Internet Exchange points, this analysis hopes to at least give a rough indication of regional characteristics.

2.4.2 Challenges and Limitations affecting the Data Analysis

There were cases where the traceroute output seemed to show "loops"; for example, AS1 followed by AS2, and then back to AS1. Further investigation revealed that in these cases, the same administrative authority managed both ASes. In cases like these, each unique AS number was considered only once. In addition, one of the small content providers (6students) had many traces that could not be determined as having reached the destination network. Results for 6students were computed based only on the traces that were complete from source to destination.

Chapter 3

Results and Analysis

Tables 4 and 5 summarize the results of the data collection, followed by an analysis of the results using the metrics described in section 2.4.1. In the analysis for each metric, the data for each group of content providers is sorted in ascending order to generate the respective graphs.

Table 4. Summary of Results for Large Content Providers

Large Content Providers	APL	NCN	AMN	AIX	GEOB
Google	3.37	37	0.173	0.08	2
Facebook	3.88	18	0.529	0.37	2
Yahoo	4.81	17	0.577	0.31	3
Baidu	4.85	1	0.615	0.02	2
Live	3.96	28	0.173	0.4	3
Wikipedia	4.35	10	0.558	0.31	2
Twitter	4.08	7	1.06	0.06	2
LinkedIn	4.43	3	0.275	0.22	2
Amazon	3.96	13	0.529	0.24	3
Wordpress	4.58	9	1.173	0.23	3
Average	4.23	14.30	0.57	0.22	2.4

Table 5. Summary of Results for Small Content Providers

Small Content Providers	**APL**	**NCN**	**AMN**	**AIX**	**GEOB**
YSU	5.04	1	0.04	0.15	2
Ampreonrecorder	4.61	4	0.96	0.2	3
Popdetective	4.35	6	0.83	0.13	3
Endlessanalog	4.36	7	1.04	0.08	3
Developphp	4.19	6	1.06	0.27	3
Italianfoodforever	4.54	5	0.86	0.2	3
Novaroma	4.24	6	1.25	0.12	2
Mixonline	4.57	4	0.81	0	2
6students	5.25	1	0.96	0.37	3
Liveidiots	4.41	1	1.48	0	3
Average	4.56	4.1	0.93	0.152	2.7

3.1 Average AS Path Length

For the large content providers, the shortest average AS path lengths were Google with 3.37, Facebook with 3.88, and Live and Amazon both with 3.96. The longest average path length was Baidu with 4.85, although Yahoo was not far behind with 4.81. Among the small content providers, the shortest AS path lengths were Developphp with 4.19, Novaroma with 4.24, Popdetective with 4.35, and Endlessanalog with 4.36. The longest average AS path lengths were 6students with 5.25, and YSU with 5.04. Figure 2 shows the comparison of average path lengths for large and small content providers. The average for large providers was 4.23 compared to 4.56 for small providers, for a difference of 0.33.

Figure 2. Average Path Length shortest to longest

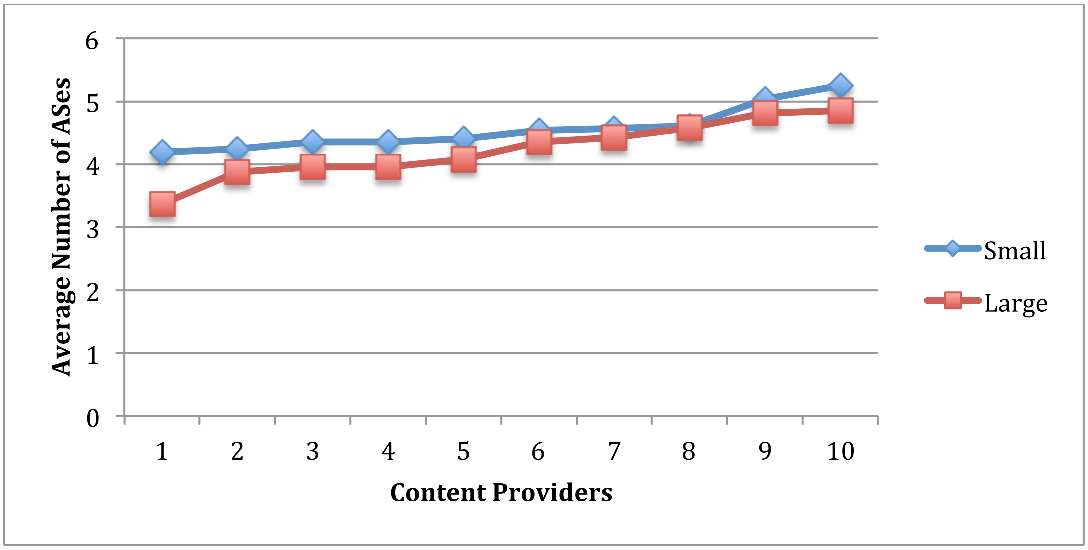

Large	APL	Small	APL
1. Google	3.37	**1.** Developphp	4.19
2. Facebook	3.88	**2.** Novaroma	4.24
3. Amazon	3.96	**3.** Popdetective	4.35
4. Live	3.96	**4.** Endlessanalog	4.36
5. Twitter	4.08	**5.** Liveidiots	4.41
6. Wikipedia	4.35	**6.** Italianfoodforever	4.54
7. Linkedin	4.43	**7.** Mixonline	4.57
8. Wordpress	4.58	**8.** Ampreon	4.61
9. Yahoo	4.81	**9.** YSU	5.04
10. Baidu	4.85	**10.** 6students	5.25

3.2 Number of networks to which a content provider is connected

Among the large content providers, the most connected by far was Google, connecting to 37 unique networks. The next most connected was Live with 28. This result is not unexpected in light of earlier studies that have shown the both Google and Microsoft appear to be deploying their own wide area networks.[22][23] LinkedIn, Twitter and Wordpress connected to three, six and seven networks respectively. LinkedIn is paired with the content delivery network Limelight (AS 22822) in 37 locations. The least connected was Baidu, connecting to only one unique network, ChinaNet (AS 4134).

For small content providers, Endlessanalog was the most connected with seven unique networks, while Popdetective, Developphp and Novaroma all connected to six. The least connected were 6students, Liveidiots and YSU, with only one connected network each.[24] Figure 3 shows the comparison between large and small content providers based on number of directly connected networks. The difference in the averages between large providers (14.30) and small providers (4.1) was 10.2.

---

[22] (Gill, Arlitt and Li)

[23] (Drivere, Ogbonna and Gundla)

[24] YSU is exclusively connected to the Columbus, Ohio-based OARnet (AS 3112). Liveidiots is exclusively connected to the major network Cogent (AS 174). 6students only connected to one network in the 20 complete traceroute paths, and no conclusion could be determined from the data regarding the other 32 paths.

Figure 3. Number of directly connected networks

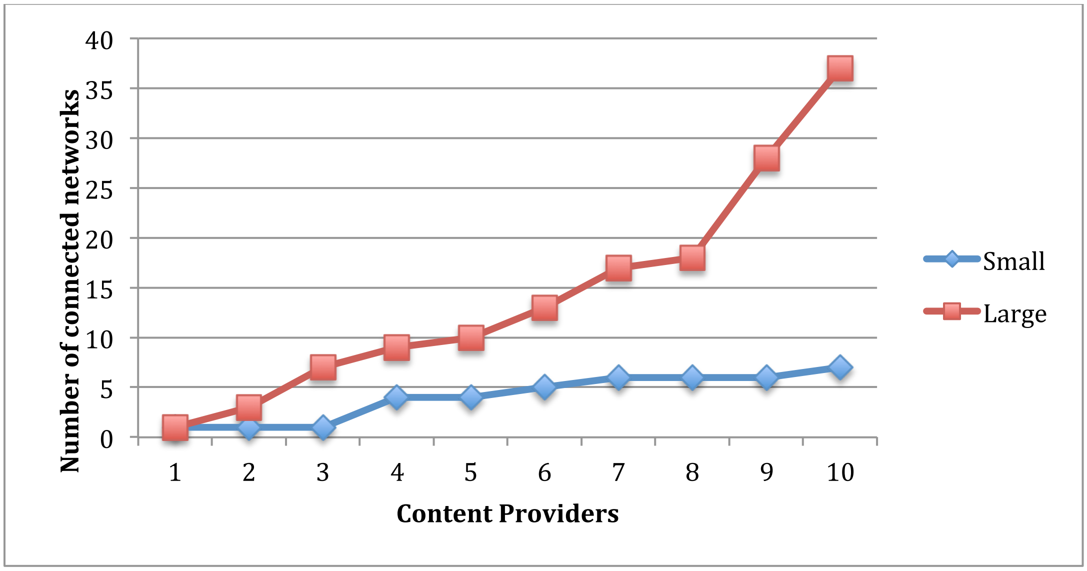

Large	NCN	Small	NCN
1. Baidu	1	**1.** YSU	1
2. Linkedin	3	**2.** 6students	1
3. Twitter	7	**3.** Liveidiots	1
4. Wordpress	9	**4.** Ampreonrecorder	4
5. Wikipedia	10	**5.** Mixonline	4
6. Amazon	13	**6.** Italianfoodforever	5
7. Yahoo	17	**7.** Popdetective	6
8. Facebook	18	**8.** Developphp	6
9. Live	28	**9.** Novaroma	6
10. Google	37	**10.** Endlessanalog	7

3.3 Average Number of Major Networks per Path

In the large content provider group, Wordpress and Twitter were the highest with 1.173 and 1.06 respectively. Google and Live were tied with the lowest at 0.173. It is interesting to note that these figures appear to show an inverse relationship to the connectivity results in Section 3.2 for Google and Live.

For the small content providers, Liveidiots was the highest at 1.48 major networks per path, with Novaroma, Developphp and Endlessanalog at 1.25, 1.06 and 1.04 respectively. The lowest was YSU with .04 major networks per path. YSU's status as an educational institution may have an effect on this result. A closer look at the AS paths in the traces issued to YSU shows a high occurrence of experimental research and academic networks like Internet 2 and National Lambda Rail. Since most of the traceroutes were issued from other colleges and universities, the YSU data may show the bias of an "educational" Internet topological vantage point. The next lowest figures for major networks per path were Mixonline, Popdetective, and Italianfoodforever, with 0.81, 0.83 and 0.86 respectively. Figure 4 shows the comparison between large and small content providers. The large provider average was 0.57, and the small provider average was 0.93 for a difference of 0.36.

Figure 4. Average number of major networks per path

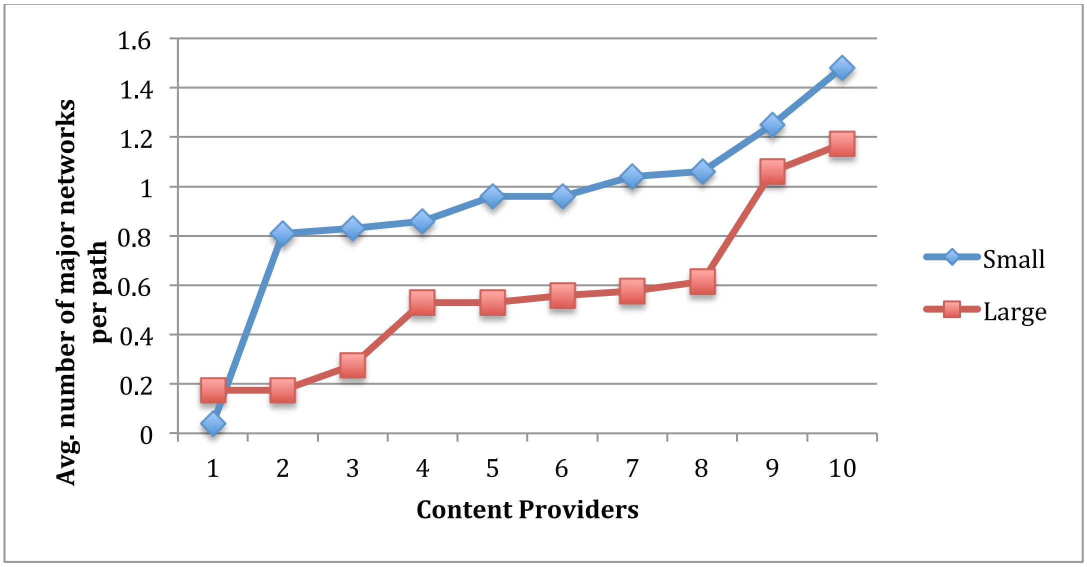

Large	AMN	Small	AMN
1. Google	0.173	**1.** YSU	0.04
2. Live	0.173	**2.** Mixonline	0.81
3. Linkedin	0.275	**3.** Popdetective	0.83
4. Facebook	0.529	**4.** Italianfoodforever	0.86
5. Amazon	0.529	**5.** Ampreon	0.96
6. Wikipedia	0.558	**6.** 6students	0.96
7. Yahoo	0.577	**7.** Endlessanalog	1.04
8. Baidu	0.615	**8.** Developphp	1.06
9. Twitter	1.06	**9.** Novaroma	1.25
10. Wordpress	1.173	**10.** Liveidiots	1.48

3.4 Average Number of Internet Exchanges per Path

Among the large content providers, Live and Facebook had the highest numbers, with 0.4 and 0.37 respectively. The lowest average number of IXs per path was found for Baidu with .02, with Twitter and Google showing .06 and .08 respectively. One interesting note about Twitter and Google: While Google was tied for lowest number of major networks per path at 0.173, Twitter was the second highest in major networks per path at 1.06, connecting directly to the major networks Global Crossing (26 times) and Level 3 (15 times).

Small content providers overall showed less of an IX presence, with the exception of 6students at 0.37. It was observed that out of 52 traceroute queries issued to 6students, only 20 were able to reach the intended destination. Out of the 20 successful queries, 17 showed at least one IX in the AS path. Among the unsuccessful queries, none showed an IX in the path. Figure 5 shows the comparison between the two groups. The averages for large and small providers were 0.22 and 0.152 respectively, for a difference of .068.

Figure 5. Average number of Internet exchanges per path lowest to highest

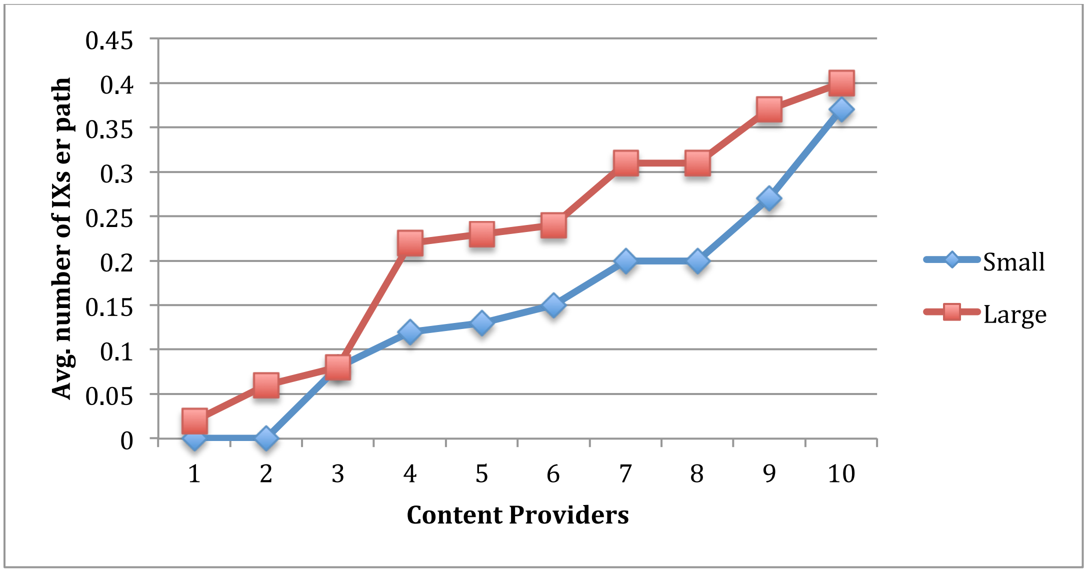

Large	**AMN**	**Small**	**AMN**
1. Google	0.173	**1.** YSU	0.04
2. Live	0.173	**2.** Mixonline	0.81
3. Linkedin	0.275	**3.** Popdetective	0.83
4. Facebook	0.529	**4.** Italianfoodforever	0.86
5. Amazon	0.529	**5.** Ampreon	0.96
6. Wikipedia	0.558	**6.** 6students	0.96
7. Yahoo	0.577	**7.** Endlessanalog	1.04
8. Baidu	0.615	**8.** Developphp	1.06
9. Twitter	1.06	**9.** Novaroma	1.25
10. Wordpress	1.173	**10.** Liveidiots	1.48

3.5 Geographical Characteristics

For the large content providers, Google, Facebook, Baidu, Wikipedia, Twitter and LinkedIn all showed only the source and destination regions in their various paths, for a maximum of two regions per path. Yahoo had one path that originated in Ottawa, Canada (ARIN) and passed through AS 9264 (APNIC) before reaching Yahoo's AS 36752 (ARIN).[25] Live had two paths with ARIN sources (Georgia and Waterloo) that passed through AS 27750 (LACNIC) before reaching Live's AS 8075 (ARIN).[26] The paths originating in Argentina (LACNIC) for Amazon and Wordpress passed through AS 6762 (RIPE) before reaching ARIN-based Amazon (AS 16509) and Wordpress (AS 13768).[27]

Among small content providers, only YSU, Novaroma and Mixonline showed no unexpected regional AS crossings. Paths originating in Brazil and Argentina (LACNIC) for Popdetective and Endlessanalog crossed through AS 6762 before reaching the destination AS 26496 (ARIN). Ampreonrecorder had paths in Brazil and Argentina that crossed through RIPE-based AS 6762 and AS 12956 respectively before arriving at AS 36476 (ARIN).[28] AS 6762 also made appearances in the Argentina paths of Developphp, Italianfoodforever, 6students and Liveidiots, all with ARIN-based AS numbers. Italianfoodforever also had AS 10026 (APNIC) in the path originating from Massachusetts.[29]

---

[25] AS 9264 is registered to Taiwan-based Academic Sinica Network.

[26] AS 27750 is registered to Cooperación Latino Americana de Redes Avanzadas, Montevideo, UY.

[27] AS 6762 is registered to Telecom Italia Sparkle S.p.A., an international carrier that has a CAIDA AS rank of 33.

[28] AS 12956 is registered to Madrid, Spain-based Telefonica Wholesale Network

[29] AS 10026 is registered to Pacnet Global Ltd and has a CAIDA AS rank of 25.

Chapter 4

Conclusions and Future Work

In this study, some interesting results were observed. While the average AS path length for small content providers was only slightly longer than the average for large content providers, more noticeable differences were evident in the other categories. Overall, small content providers showed more major networks per path, while large content providers had more Internet exchanges per path, and were directly connected to more unique networks. One of the large content providers was observed to use the services of a content delivery network.[30] Three of the small content providers and one of the large providers were each connected to only one network. The following observations represent some interesting changes in Internet topology, as seen from the content provider vantage point, from the traditionally defined Tier-1/Tier-2/Tier-3 hierarchical structure:

1. Some large content providers appear to be deploying their own wide area networks, decreasing their reliance on large Internet service providers.
2. Internet exchange points are common in the paths of both large and small content providers.

Some of the possible reasons for these changes may include greater control over content delivery, and less vulnerability to changing business relationships. The increasing use of cloud computing, multimedia applications, and software-as-a-service may be a strong incentive. If this business model is successful for very large content

---

[30] LinkedIn uses the content delivery network Limelight Networks in many geographical locations. See Section 4.2.

providers like Google and Microsoft, smaller content providers may also want to find ways to deliver their content that have been shown to be more efficient and cost effective. The increasing presence of Internet exchange points may be an indication that is happening already. The peering, transit, and business relationships between networks and IXPs is different than it is between traditionally categorized Tier-1/Tier-2/Tier-3 ISPs. If peering with an IX provides a clear advantage in terms of quality of service vs. cost, this trend can be expected to continue. It will be interesting to see how major ISPs and content providers respond to these changes. In addition, there is a possibility for future research on the security concerns that may arise for IP traffic that is transited through Internet exchange points It is hoped that the results of this study provide an answer of *yes* to the question posed in the title.

A future plan for the data collected for this study includes investigating a graph model of the topology in a test bed setting. Implementing a model of the topology in a test bed would allow researchers and students to do the following:

1. Understand possible attacks on the Internet (i.e., Distributed Denial of Service attacks)
2. Study the likelihood of a “clean slate” network design that includes all the desired characteristics BGP currently lacks
3. Develop teaching tools that can give a better insight on the structure of the Internet
4. Investigate the security vulnerabilities of the current topology and develop possible defense mechanisms

Bibliography

1. Alexa. Alexa Top 500 Sites on the Web. May 2011 <http://www.alexa.com>.

2. American Registry of Internet Numbers (ARIN). ARIN whois. October 2011 <http://www.arin.net>.

3. Amini, Lisa, Anees Shaikh and Henning Schulzrinne. "Issues with inferring Internet topological attributes." www.elsevier.com. Elsevier B.V., 8 August 2003.

4. Asia Pacific Network Information Center (APNIC). APNIC whois. October 2011 <http://www.apnic.net>.

5. BGP: the Border Gateway Protocol, Advanced Internet Routing Resources. Global Internet Exchange Points/BGP Peering Points/IXP. October 2011 <http://www.bgp4.as/internet-exchanges>.

6. —. Global Internet Exchange Points/BGP Peering Points/IXP. October 2011 <http://www.bgp4.as/looking-glasses>.

7. Case Western Reserve University. DipZoom: Deep Internet Performance Zoom. May-September 2011 <http://www.case.edu>.

8. Crovella, Mark and Balachander Krishnamurthy. Internet Measurement: Infrastructure, Traffic & Applications. West Sussex: John Wiley & Sons Ltd., 2006.

9. Drivere, Aleisa, et al. "Identifying the Impact of Content Provider Wide Area Networks on the Internet's Topology." CAHSI Annual Meeting. 2011.

10. Gill, Phillipa, et al. "The Flattening Internet Topology: Natural Evolution, Unsightly Barnacles or Contrived Collapse?" Passive Active Measurment (PAM). Berlin: Springer-Verlag, 2008. 1-10.

11. Haynal, Russ. Internet Growth Charts. October 2011 <http://www.navigators.com>.

12. Internet Partners Inc. Home Page for Whois.ipinc.net. October 2011 <http://whois.ipinc.net/index.shtml>.

13. Internet World Stats. Internet World Stats Population and Usage Statistics. May 2011 <http://www.internetworldstats.com>.

14. Kurose, James F. and Keith W. Ross. Computer Networking: A Top Down Approach. New York: Pearson Education Inc., 2010.

15. Latin American/Caribbean Network Information Center (LACNIC). LACNIC Registration Services Whois. October 2011 <http://www.lacnic.net>.

16. Luckie, Matthew, Young Hyun and Bradley Huffaker. "Traceroute Probe Method and Forward IP Path Interference." IMC '08. Vouliagmeni, Greece: ACM, 2008.

17. Peterson, Larry, et al. "Using PlanetLab for Network Research: Myths, Realities, and Best Practices." June 2005. www.measurement-lab.org. October 2011 <http://www.measurment-lab.org/files/pdn/PDN-05-028/pdn-05-028.pdf>.

18. Telia Sonera . Telia Sonera Looking Glass. October 2011 <http://lg.telia.net>.

19. The Cooperative Association for Internet Data Analysi (CAIDA). AS Rank:AS Ranking. October 2011 <http://www.caida.org>.

20. The Reseaux IP Europeens Network Coordination Center (RIPE NCC). RIPE Database Query. October 2011 <http://www.ripe.net>.

21. The University of Oregon Route Views Project. October 2011 <http://routeviews.org>.

22. Traceroute.org. Traceroute.org. October 2011 <http://traceroute.org>.