

ANALYSIS OF INTERNET TOPOLOGY DATA

by

Hao (Johnson) Chen

B.E. in Computer Science, Shen Zhen University, 1997

A THESIS SUBMITTED IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF
MASTER OF SCIENCE

In the School
of
Computing Science

© Hao (Johnson) Chen 2004

SIMON FRASER UNIVERSITY

April 2004

All rights reserved. This work may not be
reproduced in whole or in part, by photocopy
or other means, without permission of the author.

Approval

Name: Hao (Johnson) Chen

Degree: Master of Science

Title of Thesis: Analysis of Internet Topology Data

Examining Committee:

Chair: Dr. Rodney Vaughan

Senior Supervisor
Dr. Ljiljana Trajković

Supervisor
Dr. Joseph Peters

SFU Examiner
Dr. Lou Hafer

Date Approved:

Abstract

Internet topology describes the arrangement of nodes and their connecting links in the Internet. Discovering Internet topology is important for analyzing routing protocols, Internet robustness, and Internet resilience. Recent research results dealing with Internet topology, such as the discovery of power-laws and the application of spectral analysis to Internet topology data, have increased the need for more complete datasets and more rigorous interpretations.

In this thesis, we examine two AS (Autonomous System) datasets: Route Views and RIPE. They are collected from BGP (Border Gateway Protocol) routing tables and have been extensively used by the research community. These two datasets are large compared to datasets previously employed in various research studies.

Spectral approach (Laplacian analysis) enables a more in-depth analysis of data that typical approach employing adjacent matrix analysis cannot achieve. In this thesis, we create topology graphs from the two datasets, calculate the largest eigenvalues of the normalized Laplacian matrices of the graphs, and use the results to identify distinct cluster characteristics. These clusters characteristics could not be otherwise identified from the collected data.

Geographic location of Autonomous Systems (ASs) may influence inter-domain routing policies and AS connectivity. Access-providers, the owners of ASs, may prefer that incoming traffic be localized to their specific geographic areas. Most participating ASs (ASs that contribute routing tables) in Route Views and RIPE are located in North America and Europe, respectively. In order to analyze geographically related AS routing

policies for controlling incoming traffic, we propose a notion of “reverse pairs.” Our analysis shows that the effect of routing policies is not negligible.

Dedication

My thesis, degrees, and achievements could not be completed successfully without the endless and selfless support from beloved ones: my parents, Cai Tian (蔡甜) and Guangshen Chen (陈广深), my sister Feihong Chen (陈飞虹), and my girlfriend Fei Chen (陈菲). I love you all dearly!

Acknowledgements

My senior supervisor, Prof. Ljiljana Trajković has always stimulated me to do better work. From her, I learned so much both intellectually and personally. I am lucky that I have a good special friend like her during my study in SFU.

Many of thanks to members of the School of Computing Science: Joseph Peters who was my research and study mentor, Lou Hafer who provided tremendous help with the ACM contest simulation work and when I was TA for CMPT 471, Kersti Jaager, Tony Dixon, Elizabeth Zook, Brad Bart, Jens Happe, and Jeff Orchard. They have helped me a great deal with my graduate studies.

Many thanks to members of Networking Communications Laboratory in SFU: Renju Narayanan, Nenad Lasković, and Andre Dufour who assiduously proofread my thesis, Tony Feng and Vladimir Vukadinović who stayed late nights with me, and other members with whom I have shared many fun moments.

I spent a wonderful time at SFU over the past three years. It was a memorable experience in my life.

Table of Contents

Approval	ii
Abstract	iii
Dedication	v
Acknowledgements	vi
Table of Contents	vii
List of Figures	ix
List of Tables	x
List of Abbreviations and Acronyms	xi
Chapter 1 Introduction	1
Chapter 2 Internet Topology	5
2.1 Topology hierarchy of the Internet	5
2.2 Internet routing	6
2.2.1 Internet Protocol	6
2.2.2 Classless Inter-Domain Routing (CIDR)	7
2.2.3 Autonomous Systems (ASs)	8
2.2.4 Open Shortest Path First (OSPF) Protocol	9
2.2.5 Border Gateway Protocol (BGP)	10
2.3 Internet topology data on the AS level	12
2.3.1 Sources of BGP routing tables	13
2.3.2 Sources of <i>traceroute</i> data	14
2.3.3 BGP data vs. <i>traceroute</i> data	15
2.3.4 Limitations of the used Internet topology data	15
2.4 Power-law distributions	19
2.4.1 Power laws in Internet topology	19
2.4.2 Origin of power laws in Internet topology	21
2.5 AS relationships	22
2.6 Geographic properties of the Internet topology	26
Chapter 3 Spectral Graph Theory	28
3.1 Spectra of a graph	28
3.2 Eigen-analysis	35
Chapter 4 Analysis of Internet Data	37
4.1 Observations of Internet topology datasets	37
4.2 Spectral analysis of the AS Internet topology	38
4.3 Reverse pairs	43
Chapter 5 Conclusions	49
Appendix A List of participating ASs in the Route Views and RIPE datasets	51
Appendix A-1 Participating ASs in Route Views	51

Appendix A-2 RRC information in RIPE	53
RRC00 - RIPE NCC Peer List:.....	53
RRC01 Peer List:.....	54
Appendix B List of programs and script used to preprocess data	56
Appendix B-1 Data filtering procedure.....	56
Appendix B-2 Program (Exact.java) used to extract AS_Path from data.	56
Appendix B-3 Program (AdjMatrix.java) used to build adjacent matrix.....	57
Appendix B-4 Program (Deg.java) used to calculate degrees of ASs.	59
Appendix B-5 Program (gen_l.m) used to generate the Normalized Laplacian matrix.	59
References	61

List of Figures

Figure 2.1 A simple physical topology of the Internet.	6
Figure 2.2 Four classes of IP addresses.....	7
Figure 2.3 An example of a selected announcement prefix.	26
Figure 3.1 An example of a multigraph.	28
Figure 3.2 An example of a complete graph.	29
Figure 3.3 An example of a bipartite graph.	29
Figure 3.4 An example of a clique graph.	30
Figure 3.5 Adjacency matrix and Laplacian matrix of a sample graph.....	31
Figure 4.1 AS degree distribution in Route Views and RIPE datasets.....	39
Figure 4.2 An example demonstrates how characteristics valuation process works.	40
Figure 4.3 Spectral views of AS connectivity in two datasets.	41
Figure 4.4 An example of reverse pair path.	48

List of Tables

Table 2.1 A snippet of a BGP routing table.....	11
Table 2.2 Comparison of data sources.	12
Table 3.1 Spectrum and eigenvectors of a sample graph.....	31
Table 4.1 General statistics of the Route Views and RIPE datasets.....	37
Table 4.2 Assigned numbers of twenty ASs with the largest node degrees.	38
Table 4.3 An example of a small cluster in the RIPE dataset.	42
Table 4.4 ASs (they belong to a cluster in RIPE dataset) in Route Views dataset are separated into two clusters (classified by element values).	43
Table 4.5 ASs (with total degree larger or equal to 10) among reverse pairs in RIPE (a), Route Views (b).	46
Table 4.6 General statistics of routes built from reverse pairs.	47
Table 4.7 Routes built from reverse pairs with the maximum number of hops equal to 3.....	47

List of Abbreviations and Acronyms

AS	Autonomous System
ASN	Autonomous System Number
BGP	Border Gateway Protocol
CAIDA	The Cooperative Association for Internet Data Analysis
CIDR	Classless Inter-Domain Routing
DNS	Domain Name System
HOT	Highly Optimized Tolerance
IANA	Internet Assigned Numbers Authority
ICANN	Internet Corporation for Assigned Names and Numbers
ICMP	Internet Control Message Protocol
IP	Internet Protocol
IRR	Internet Routing Registry
ISP	Internet Service Provider
IX	Exchange Point
NIC	Network Information Center
OSPF	Open Shortest Path First Protocol
POP	Point of Presence
RIPE	Réseaux IP Européens
RIR	Regional Internet Registry
RRC	Remote Route Collectors
SA Prefix	Selective Announced Prefix
TCP/IP	Transmission Control Protocol/Internet Protocol

Chapter 1 Introduction

In spite of Internet's exponential growth, certain characteristics of its topology remain unchanged. Better understanding of these invariants may contribute to future Internet research and development, such as new protocol designs. Internet topology describes the arrangement of network nodes and their connecting links. There are two types of topology: physical and logical. Physical topology specifies the physical layout of network nodes and links. Logical topology, the topic of this thesis, refers to the paths of traffic flowing from node to node.

Because of the large size of Internet, researchers frequently restrict analysis of its topology to the autonomous system (AS) level. AS is a single network or a group of networks that have a coherent routing policy. Instead of dealing with all hosts and routers, researchers can consider only 15,000 active ASs [2] to analyze Internet topology. Unless otherwise specified, Internet topology and Internet routing discussed in the thesis are on AS level.

When Internet topology information is reduced to the AS level, complete AS data becomes of great importance for the validity of research results. Unfortunately, due to the underlying complex mechanisms, it is impossible to acquire complete Internet AS data. In general, studies of Internet topology rely on limited Internet AS data.

Internet AS data either emanate from the Border Gateway Protocol (BGP) routing tables or are inferred from IP addresses collected by *traceroute* (a TCP/IP utility). Data from BGP routing tables of the Route Views project from the University of Oregon [31] and data from routing tables of Réseaux IP Européens (RIPE) [29] have been extensively used by the research community [9], [15], [26], [37]. Of many important

properties of Internet routing, interconnectivity status and geographic location of ASs are commonly studied to analyze Internet topology.

The Internet connects thousands of ASs operated by many distinct administrative domains, such as Internet Service Providers (ISPs), companies, and universities. Border Gateway Protocol (BGP) is often used in inter-AS routing. A key feature of BGP is to allow ASs to choose their own administrative policy when ASs are either selecting the best route, announcing, or accepting routes. Commercial agreements between pairs of administrative domains play an important role in determining routing policies. These agreements determine relationships between ASs. In general, AS relationships can be classified into customer-provider, peering, mutual-transit, and mutual-backup agreements [21], [22]. AS relationships is an important factor in shaping the structure of the Internet and in influencing the Internet's end-to-end performance characteristics.

The geographic location of ASs provides insight into the structure and functionality of the Internet. Points of presence in the Internet are physical access locations where ISPs connect to the Internet. Chang et al., [10] observed that the number of points of presence may be a controlling force in determining Internet topology. Furthermore, geographic and network locations of the end-hosts, which are residing within ASs, may influence the end-to-end performance of routes [33].

Discoveries from collected Internet topology data are sometimes contradicting, which draws people's attention to what source of topology data to be used. In 1999, Faloutsos et al., [15] discovered power-laws in the degree distributions of Internet graphs. They constructed graphs using Internet topology data from Route Views. Later, Chang et al., [9] observed the Weibull degree distribution of Internet graphs. Data used to construct the graphs were derived from Route Views, IRR information of RIPE, 11 public route servers, and several Looking Glass sites (sites where their BGP summary

information can be accessed by the public). Chang et al., also suggested that data from Route Views may be incomplete: after incorporating the RIPE dataset into the Route Views dataset, they found that the combined dataset had $\sim 40\%$ more AS connections and 2% more ASs than those found in the Route Views dataset.

In order to analyze the immense volume of the Internet topology data, it is crucial to apply the appropriate methods. In the thesis, we choose spectral methods. Eigenvalues of a graph are closely related to many basic topological properties, such as the diameter, the number of edges, and the numbers of spanning trees. Power-law held for relations between the eigenvalues of Internet graphs and their ranks has been identified in [15]. Vukadinović et al., [37] used the spectrum (set of eigenvalues) to distinguish actual AS interconnection graphs from synthetic graphs. Clustering characteristics have also been studied [26] by employing an approach similar to the use of eigenvectors in the works of Fiedler [16].

In order to address concerns about Internet topology data raised by Chang et al., [9] and to further analyze characteristics of Internet topology, in this thesis we examine datasets from the Route Views project and RIPE. We choose these two data sources because they have been widely used in the research community [9], [15], [26], [37] and provide the most complete available Internet topology data. We first use spectral methods to analyze the two datasets and show how eigenvectors corresponding to the second smallest and the largest eigenvalues can partition data and indicate clusters in the Internet graph, respectively. Because of the distinguished geographic properties of the Route Views and RIPE datasets, we also propose the notion of “reverse pairs” to study Internet routing policies within the two datasets. This study is unique because it employs both datasets together to analyze geographically related Internet routing

policies. Our results suggest that locating “reverse pairs” in the Route Views and RIPE datasets may assist in analyzing routing policies on incoming traffic in the Internet.

The thesis is organized as follows. In Chapter 2, we provide background on Internet routing and descriptions of main topology datasets, power-law distribution in Internet topology, AS relationships, and geographic properties of Internet topology. In Chapter 3, we introduce spectral graph theory. Our results are presented in Chapter 4. We conclude with Chapter 5.

Chapter 2 Internet Topology

In this chapter, we introduce the background analysis of Internet topology data. We address topology hierarchy of the Internet, Internet routing, Internet topology data, power laws of Internet topology, as well as AS relationships and geography. AS relationships and geography are two important Internet routing properties that are related to Internet topology.

2.1 Topology hierarchy of the Internet

The Internet can be perceived as a multi-tiered architecture consisting of Internet service providers (ISPs). There are thousands of ISPs that provide Internet access to individual users and companies. Tier-1 ISPs (backbone ISPs) are a handful of ISPs that maintain routing tables. These routing tables are default-free because they contain complete reachability information of all globally network-layer addresses reachable throughout the Internet. Tier-1 ISPs operate extensive high-speed backbone networks. Most other ISPs (regional ISPs or tier-2, 3, ..., ISPs) derive their connectivity from larger ISPs. Many backbone ISPs interconnect with each other at Internet exchange points (IXes). A simple physical topology of the Internet is shown in Figure 2.1.




Figure 2.1 A simple physical Internet topology that has Exchange Points, Backbone ISPs and Regional ISPs.

2.2 Internet routing

Data is routed from one computer to another in the Internet. Internet routing consists of a variety of protocols and techniques.

2.2.1 Internet Protocol

The Internet Protocol (IP) is a protocol used to send data between hosts in the Internet. Each host has at least one IP address that uniquely identifies it among all other computers within the Internet. IP defines how the data will be divided into packets. Each packet contains an IP address. These packets travel across the Internet by different routes and arrive at the destination in a varying order.

The original Internet Protocol (Internet Protocol version 4 or IPv4) divides IP addresses into four classes of address structure: Classes A through D. They are shown in Figure 2.2. An IP address is a 32-bit number that identifies each sender or receiver of packets sent across the Internet. Each of these four classes (A-D) allocates one portion

of the 32-bit Internet address format to a network address. The remaining portion of the 32-bit address is allocated to the specific host machines within the specified network.

Class A

0	Network (7 bits)	Local address (24 bits)
---	------------------	-------------------------

Class B

10	Network (14 bits)	Local address (16 bits)
----	-------------------	-------------------------

Class C

110	Network (21 bits)	Local address (8 bits)
-----	-------------------	------------------------

Class D

1110	Multicast address (28 bits)
------	-----------------------------

Figure 2.2 Four classes of IP Address.

2.2.2 Classless Inter-Domain Routing (CIDR)

One of the most commonly used address classes is Class B. It allocates space for up to 65,533 host addresses. A company requires a block of Class B address if it needs more than 254 host machines, even what actually it needs are far fewer than 65,533 host addresses. This allocation would "waste" most of the block of allocated addresses. For this reason, the Internet was running out of address space more quickly than necessary. Classless Inter-Domain Routing (CIDR) solved the problem by providing a new and more flexible way to specify network addresses in routers. CIDR requests that each IP address has a *network prefix* that identifies either an aggregation of network gateways or an individual gateway. A gateway is a network point that acts as an entrance to another network. The length of the network prefix is specified as part of the IP address. The length varies depending on the number of bits needed. A destination IP address or a route that describes many possible destinations has a shorter prefix and is said to be less specific. A longer prefix describes a destination gateway more specifically. Routers are required to use the most specific or the longest network prefix in the routing table when forwarding packets.

An example of a CIDR network address is: 192.168.1.0/18. The "192.168.1.0" is the network address. The entry "18" identifies that the first 18 bits are the network part of the address. This leaves the last 14 bits for specific host addresses. CIDR allows one routing table entry represent an aggregation of networks that exist in the forward path.

2.2.3 Autonomous Systems (ASs)

An autonomous system can be either a single network or a group of networks controlled by a common network administrator (or group of administrators) on behalf of a single administrative entity. An administrative entity can be a university, a business enterprise, or a business division. An AS is a connected group of one or more IP prefixes run by one or more network operators that have a defined routing policy. Routing policy is defined as the set of routing decisions. The exchange of routing information between ASs is subject to routing policies. An autonomous system is assigned a globally unique number, called an Autonomous System Number (ASN).

An AS usually maintains a range of IP addresses. Users are assigned IP addresses by ISPs. ISPs usually obtain allocations of IP addresses from their respective Regional Internet Registry (RIR):

- APNIC (Asia Pacific Network Information Centre)
- ARIN (American Registry for Internet Numbers)
- LACNIC (Regional Latin-American and Caribbean IP Address Registry)
- RIPE NCC (Réseaux IP Européens)
- AfriNIC (an African Regional Internet Registry).

IANA (Internet Assigned Numbers Authority), which is under the control of ICANN (Internet Corporation for Assigned Names and Numbers), delegates local

registrations of IP addresses to RIR. As seen from the RIR list above, each RIR allocates addresses for a specific geographical area.

2.2.4 Open Shortest Path First (OSPF) Protocol

The Internet consists of a set of ASs, each having a set of gateways. A gateway is a network point that acts as an entrance to another network. An Interior Gateway Protocol (IGP) is a protocol for exchanging routing information between gateways within an AS, while Exterior Gateway Protocol (EGP) serves for exchanging routing information between two neighbor gateway hosts in a network of ASs. OSPF and BGP are more recent versions of IGP and EGP, respectively.

OSPF allows collections of adjacent networks and hosts to be grouped together. Such a group, together with the routers having interfaces to any of the included networks, is called an area. A copy of the link-state routing algorithm is run at each area. The topology of an area is invisible from the outside of the area. Routers internal to a given area do not know the detailed topology external to the area. This isolation of knowledge enables the protocol to reduce routing traffic substantially, as compared to treating the entire AS a single link-state domain. The backbone of the AS consists of networks not contained in any area, routers attached to these networks and routers that belong to multiple areas. The backbone must be contiguous.

When an AS is split into OSPF areas, routers can be divided into four overlapping categories according to router' functions:

- Internal routers are routers with directly connected networks belonging to the same area. A single copy of the link-state routing algorithm is run at an internal router.

- Area border routers are routers that attach to multiple areas. Multiple routing algorithm copies, one copy for each attached area and an additional copy for the backbone, are run at an area border router.
- Backbone routers are routers that have an interface to the backbone. Backbone routers can be area border routers or internal routers.
- AS boundary routers are routers that exchange routing information with routers belonging to other ASs. An AS boundary router has AS external routes that are advertised throughout the AS. Every router in the AS knows the path to each AS boundary router.

2.2.5 Border Gateway Protocol (BGP)

BGP is an inter-AS routing protocol. The primary function of a BGP routing system is to exchange network reachability information with other BGP systems. The network reachability information includes the list of ASs that the packet containing the reachability information traverses. This information is sufficient to prune propagation routing loops.

BGP routing information is stored in routing tables. A routing table contains a list of known routers, the addresses they can reach, and a cost *metric* associated with the path to each router. A portion of a BGP routing table example from a Cisco router is shown in Table 2.1. Each row represents a route, except for the first 2 rows. The first column contains the Status codes, which indicate one of the 6 possible statuses of a route. The last column ends with the value of the BGP Origin codes, which indicate where the route originates from. The values of all the Status codes and Origin codes are listed in the first two rows in Table 2.1.

Table 2.1 A snippet of BGP routing table.

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal						
Origin codes: i - IGP, e - EGP, ? – incomplete						
	Network	Next Hop	Metric	LocPrf	Weight	Path
*	3.0.0.0	193.251.128.22			0	5511 7018 80 i
*		202.249.2.86			0	7500 2518 1239 7018 80 i
*		217.75.96.60			0	16150 8434 3549 7018 80 i
*		216.140.2.62	6788		0	6395 7018 80 i

The remaining six columns of Table 2.1 show the following route information:

- Network: the prefix for this route. It includes a prefix length (or mask) unless it has a "classical" (pre-CIDR) length of 0, 8, 16, or 24 bits corresponding to a default route (0 bits) or a class A, B, or C. In other cases, the field is empty indicating it is another route for the prefix that appeared last. For example, in Table 2.1, the last 3 rows all correspond to routes of prefix "3.0.0.0".
- Next Hop: the BGP NEXT_HOP attribute. The next hop is the address of the AS boundary router to which traffic for this prefix will be forwarded, i.e., the next AS boundary router in the path to its destination. The address may be 0.0.0.0, indicating that the next hop is directly connected to the destination.
- Metric: the BGP MULTI_EXIT_DISCRIMINATOR attribute. It is used as a suggestion to an external AS regarding the preferred route into the AS that is advertising the metric. In the BGP route selection process, the lowest value is preferred. Its upper bound is $2^{32} - 1$.
- LocPrf: the BGP LOCAL_PREFERENCE attribute. This is an administrative preference value. In the BGP route selection process, the highest value is preferred.
- Weight: an administrative preference particular to Cisco routers. Most vendors have such a value. It is a local value, not exchanged between peers in BGP. In the BGP

route selection process, the highest value is preferred. It is often used to prefer routes for directly connected prefixes to any other routes.

- Path: the BGP AS_PATH attribute. The attribute records the ASs through which a route has been exchanged before the route was received by the router. If the field is empty, the local AS generates the route. In the BGP route selection process, a shorter AS path (fewer ASs) is preferred.

BGP hosts communicate via TCP and send updated routing table information only when a host has detected a change. Only the affected segment of the routing table is sent.

2.3 Internet topology data on the AS level

AS connectivity data may be categorized into: data derived from BGP routing tables and data derived from *traceroute* (a TCP/IP utility) paths. The research community has extensively used two sources of routing tables: data from the Route Views project from University of Oregon [31] and data from Réseaux IP Européens (RIPE) Network Coordination Center [29]. The Cooperative Association for Internet Data Analysis (CAIDA) [35] uses a *traceroute*-like approach to collect Internet topology data for further analysis.

Prior use of the two datasets from Route Views [31] and RIPE [29] is partially summarized in Table 2.2.

Table 2.2 Comparison of data sources.

	Route Views	RIPE
Faloutsos et al., [15]	Yes	No
Chang et al., [9]	Yes	No
Vukadinović et al., [37]	Yes	No
Mihail et al., [26]	Yes	Yes

Even though both datasets were used in past studies [26], they were analyzed separately. In this thesis, we analyze the new dataset obtained by combining data from Route Views and RIPE datasets.

2.3.1 Sources of BGP routing tables

The two largest available BGP routing table sources are data collected from the Route Views project [31] and from RIPE [29].

The Route Views project was originally conceived as a tool that would allow Internet operators to obtain real-time information about the global routing system from several backbone routers around the Internet. The project has currently evolved to include dozens of participating ASs and to provide Internet topology views of unicast and multicast paths in IPv4. Interesting uses of Route Views data are the AS path visualization project [1] in the National Laboratory for Applied Network Research (NLNR) [27] and the study of the IPv4 address space utilization [4]. Route Views data was also used to map IP addresses to ASs where the IP addresses originate from. CAIDA [35] has used the Route Views data together with the NetGeo [28] database to generate geographic locations of Internet hosts.

In Route Views project, a script runs every two hours to collect full BGP routes from participating BGP routers. The script connects to route-views.oregon-ix.net (the Route Views routers, and not participating BGP routers), runs “show ip bgp”, and archives the output. The Route Views project is located in Oregon, USA. Most participating ASs are in North America. The list of participating AS peers in the route-views.oregon-ix.net is given in Appendix A-1.

In contrast to the centralized way of collecting routing data in Route Views, RIPE applies a distributed approach to the data collection. The RIPE Network Coordination

Center is one of the five existing RIRs in the world today. It provides allocation and registration services primarily for the Internet users in Europe, the Middle East, and North Africa. RIPE collects routing information by using Remote Route Collectors (RRC). An RRC is a daemon running to collect default-free BGP routing information. Several RRCs have been deployed in Europe, North America and Asia. Each day, RRCs in each location collect the entire routing tables every eight hours. The collected raw data is transferred via an incremental file transfer utility *rsync* to a central storage area at the RIPE center in Amsterdam. Appendix A-2 shows the Remote Route Collectors and participating ASs. Most participating ASs reside in Europe.

2.3.2 Sources of *traceroute* data

Traceroute is a TCP/IP utility that helps to determine the route that packets take to reach a particular host. The *traceroute* utility works by increasing the “time to live” (TTL) value of each successively sent packet. The first packet sent has a TTL value of one, the second packet has a value of two, and so on. When a packet passes through a router, the router decreases the TTL value by one and forwards the packet to the next hop. When a packet with TTL of one arrives at a router, the router discards the packet and sends an Internet Control Message Protocol (ICMP) time-exceeded type packet to the sender. A router returns a time-exceeded packet to the sender if the lifetime of the sent packet expires. The *traceroute* utility uses these returning packets to produce a list of routers that the packets have traversed on the way to the destination host. The *traceroute* cannot discover AS boundaries because it is not able to access BGP routing tables in routers.

2.3.3 BGP data vs. *traceroute* data

Approaches based on *traceroute* rely on ICMP packets to determine routes to the destination. Hence, the performance of these approaches is limited. For example, many routers will filter out ICMP packets due to security reasons. Unlike approaches using BGP routing tables, *Traceroute* approaches are not able to always collect all the routes contained in routing tables. Nevertheless, *traceroute* approaches are able to collect Internet topology information that sometimes BGP approaches cannot. For instance, AS tables constructed by *traceroute* can collect more IX addresses than BGP routing tables. IXes are peering points where network providers interconnect and exchange traffic destined for each other's customers. They are a collection of border routers. In general, IXes are not ASs and not all IXes have AS number registered to them. Even if an IX has an AS number, it rarely appear in BGP paths because usually routing algorithms are not run at routers in an IX and an IX is not involved in the route selection. Therefore, BGP is blind to these interconnecting points among ASs. However, *traceroute* can directly select routers in IX ASs as destinations, and hence, IX ASs occur more often in *traceroute* AS tables than in BGP routing tables.

2.3.4 Limitations of the used Internet topology data

An important issue requires attention when using data from BGP routing tables to study the Internet's AS connectivity structure. As a policy-based protocol, BGP is not related to physical connectivity at the AS level. Rather, BGP deals with the logical AS peering relationships. Hence, BGP-derived AS connectivity may yield an incomplete picture of the physical Internet connectivity. We now address several major concerns related to the completeness of BGP-derived AS graphs.

2.3.4.1 Other available BGP routing data

Besides Route Views and RIPE, other sources of BGP routing data are available, for example, BGP data from Looking Glass project [24]. The Looking Glass sites support public troubleshooting services of their own and provide limited public access to their selected BGP routers. Some Looking Glass sites also make available BGP summary information (via the “show ip bgp summary” command). The summary contains a set of peers and aggregated BGP activity statistics, from which one can derive a set of ASs neighboring to the local ASs.

The primary purpose of the Looking Glass sites is to troubleshoot specific routing problems by query requests. In every request returned, the sites only provide a partial view of routing tables pertaining to certain arguments requested from the user (via the “show ip bgp [argument]” command). Full routing table dumps (via the “show ip bgp” command) are usually not allowed. Furthermore, Looking Glass sites are supposed to be used interactively. Substantial effort is required to obtain reasonably complete routing tables. Therefore, in Internet topology research, data from Looking Glass sites are less frequently used than data from Route Views and RIPE sites.

2.3.4.2 AS_PATH aggregation

BGP routers are allowed to compress the data in their routing tables in order to reduce the amount of required storage and the size of routing messages to be exchanged. A popular technique to reduce routing table entries is route aggregation [5]. Route aggregation may make it impossible to correctly derive AS connectivity from BGP routing tables. Route aggregation replaces a set of route specifications with a single, aggregated route specification. For example, the four prefixes:

192.168.0.0/24

192.168.1.0/24

192.168.2.0/24

192.168.3.0/24

may be replaced by :

192.168.0.0/22.

Aggregation of prefixes requires aggregation of the associated AS_PATHs. Routes are advertised between a pair of BGP routers via UPDATE messages. AS_PATH is a mandatory attribute of the UPDATE messages. It is composed of a sequence of AS path segments. Two types of AS path segment exist: AS_SET and AS_SEQUENCE. AS_SET is an unordered set of ASs a route in the UPDATE message has traversed. AS_SEQUENCE is an ordered set of ASs that a route in the UPDATE message has traversed.

AS_SETs are used in the route aggregation. AS_SETs reduces the size of the AS_PATH information by listing each AS number only once, regardless of how many times it may have appeared in multiple AS_PATHs that were aggregated. In the AS_PATH representation, AS_SETs are delimited by "{" and "}" brackets in Cisco routers or by "[" and "]" brackets in Juniper routers. For example, the following two AS_PATHs:

62011 62006 62043

62011 62007

can be aggregated into

62011 {62006, 62043, 62007}.

As a result, route aggregation can lead to information loss. In general, it is impossible to derive the original AS_PATHs from an aggregated AS_PATH.

Route aggregation may not be actively performed in the Internet presently. Internet BGP statistics [6] show that routing tables in the Internet contain only a small number of aggregated routes. The reason is or may partially be that route aggregation cannot be used to aggregate large number of routes. As the given aggregation example indicates, the aggregate summarizes individual routes. Hence, any changes in the individual route will cause the aggregate to be updated. For example, if AS 62043 is not accessible, the path information of the aggregate changes from {62006, 62043, 62007} to {62006, 62007} and the aggregate is updated. If the aggregate summarizes many routes, it may constantly flap if the routes forming the aggregate change.

We located 1,549 (~0.025%) and 1,717 (~0.022%) aggregated route entries of the entire route entries, in RIPE and Route Views datasets, respectively. Only AS pairs (two adjacent ASs in route entries) are used to construct Internet topology graphs in the thesis. We located 13 AS pairs in aggregated route entries from both RIPE and Route Views datasets. These aggregated AS pairs are erroneous data and cannot be used to derive AS graphs. Several observations of the AS pairs were made. The 13 AS pairs can be classified into two types: for the first type, one AS longs to AS_SEQUENCE while the other to AS_SET; for the second type, both ASs belong to AS_SET. We also observed that randomly-selected reserved ASNs, such as 65001, 65002, were stuffed in AS_SET in several of the 13 AS pairs in order to artificially increase the length of AS_PATH, and thus make the specific route less preferable.

We estimate the possible effect of AS aggregation on the AS graphs that we constructed. A degree of an AS is the number of connections the AS has in the AS graph. The total degree of the ASs in the 13 AS pairs we located is ~ 0.51% and 0.65% (225 and 180) of the sum of all ASs in the Route Views and RIPE, respectively. Hence, the effect can be neglected.

2.4 Power-law distributions

Faloutsos et al., [15] examined the properties of Internet topology on AS level and discovered four simple power-laws after analyzing three samples of Internet topology data.

2.4.1 Power laws in Internet topology

Modeling Internet topology is still an open research problem. Before discovering power-laws [15], several graph generators have been proposed [7], [13], and [39]. Nevertheless, the problem of generating realistic topologies is not yet solved. The selection of parameter values is often left to the intuition and the experience of researchers. Furthermore, metrics or properties such as the average outdegree (the number of outgoing connections of a node), fail to quantify topological properties and to concisely describe data distributions in the Internet topology.

The primary contribution of Faloutsos et al., [15] was to identify the existence of power-laws for the Internet topology emanating from data collected in 1998. Power-laws are expression of the form $y = \alpha x^k$, where the proportion α , and the exponent of the power law k , are constants, while x and y are measures of interest.

The following power-law is observed in examining the outdegrees of nodes:

$$d_v \propto r_v^R, \quad (2.1)$$

where d is the outdegree of a node, r_v is the rank of the node v and R is a constant. Outdegree is the number of outgoing connections from a node. Nodes are sorted in decreasing outdegree sequence. The rank of a node v is its index in the sequence.

The second power-law describes the distribution of the outdegree of the graphs by a single number α :

$$f_d \propto d^o, \quad (2.2)$$

where f_d is the frequency, d is the outdegree and o is a constant. Frequency f_d of an outdegree d is the number of nodes with outdegree d .

The third power-law quantifies the connectivity and distances between Internet nodes as:

$$P(h) \propto h^\eta, \quad (2.3)$$

where $P(h)$ is the total number of pairs of nodes within h hops, and η is a constant.

Finally, the eigenvalues λ of Internet graphs are identified:

$$\lambda_i \propto i^\varepsilon, \quad (2.4)$$

where λ_i is the eigenvalues of a graph, i is the order and ε is a constant.

Power-laws have a number of practical applications. Observations show that most Internet metrics typically follow a power-law. Exponents of power-laws capture the properties by a single number.

Chang et al., [9] questioned the completeness of data used in [15] because the analysis relied on the BGP data obtained exclusively from Route Views [31]. They suggested the inclusion of additional data. They contended that by strictly being BGP-based, the data in [15] leads to a rather incomplete picture of Internet connectivity on the AS level. The AS connectivity graphs constructed from these data typically have at least 20% fewer links than links in AS graphs constructed using “extended” source. The “extended” data that Chang et al., used consists of data from Route Views, IRR information of RIPE [29], 11 public route servers, and several Looking Glass sites. The majority of the “extended” data was obtained from Route Views and RIPE. Chang et al., showed that the connectivity-based dynamics assumed in [3] were invalid. They also

arrived at a conclusion that departed from the original power-laws discovered in Internet topology [15]. The authors performed a detailed analysis and found that, while the degree distributions resulting from the Route Views AS graphs are consistent with power-law distributions, the distribution of the “extended” AS graphs are consistent with heavy-tailed distributions, such as the Weibull distribution. Weibull distribution is a continuous probability distribution with the probability density function

$$f(x) = (k/\lambda)(x/\lambda)^{(k-1)} e^{-(x/\lambda)^k} \quad \text{for } x > 0, \quad (2.5)$$

where $k > 0$ is the shape parameter and $\lambda > 0$ is the scale parameter. The outcome of the examination of “extended” dataset raised the concern of the source of data used in Internet topology study, and indicate the need to include Internet topology data from multiple sources.

2.4.2 Origin of power laws in Internet topology

Faloutsos et al., [15] provide evidence for the existence of the four power laws in Internet topologies. They did not address the cause of the power laws existence. Barabási and Albert [3] suggest incremental growth and preferential connectivity as two possible causes for power law distributions of outdegree in a network topology. Incremental growth implies that networks are formed by continual addition of new nodes. Hence, the size of the network gradually increases. Preferential connectivity means that a new node is more likely to be connected to existing nodes that are highly connected or popular than to nodes that are less connected or less popular. Medina et al., [25] provide two additional possible causes for the existence of power laws in Internet topologies. The first possible cause describes the space distribution of a network. They conjecture that unlike random models Internet topologies have a high degree of clustering. Hence, models that generate topologies that have nodes distributed according to a heavy-tailed

distribution appear more realistic. Another possible cause for the existence of power laws is the locality of edge connection, where a new node tends to connect to the existing nodes that are close by in terms of distance.

By modifying a Highly Optimized Tolerance (HOT), a construction proposed by Carlson and Doyle [8], [14], Chang et al., [10] examined the forces that shape Internet connectivity on the AS level. HOT is a mechanism that generates power law distribution. The mechanism is motivated by biological organisms and advanced engineering technologies that are optimized either through natural selection or through engineering design, to provide robust performance despite uncertain environments. HOT suggests that power laws in these systems are due to tradeoffs between yield, cost of resources, and tolerance to risks. Carlson and Doyle show that features of HOT systems include:

- High efficiency, performance, and robustness to designed-for uncertainties (e.g., the Internet is a designed system that exhibits substantial uncertainty in the user-created environment as well as the network itself);
- Hyper-sensitivity to design flaws and unanticipated perturbations;
- Non-generic, specialized, structured configurations;
- Power laws.

Results of the experiments with the modified HOT model [10] confirm a previously reported conjecture that “the highly variable degree distribution may arise merely from its correlation with a highly variable size distribution” [34].

2.5 AS relationships

An AS is a group of connected networks administrated by one or more network operators. Each AS has a consistent routing policy defined as a set of routing decisions.

The exchange of routing information between ASs is subject to routing policies. Consider the case of two ASs, A and B, exchanging the routing information:

$$\text{NET1} \dots \text{AS_A} \Leftrightarrow \text{AS_B} \dots \text{NET2} .$$

AS_A knows how to reach a network with prefix NET1. It is irrelevant whether NET1 belongs to AS_A or to an AS that exchanges routing information with AS_A, either directly or indirectly. We only assume that AS_A knows how to direct packets to NET1. Likewise, AS_B knows how to reach NET2. In order for traffic to flow from NET2 to NET1 between AS_A and AS_B, AS_A has to announce NET1 to AS_B using an exterior routing protocol. This implies that AS_A is willing to accept traffic from AS_B directed to NET1. Routing policy comes into play when AS_A decides to announce NET1 to AS_B. For traffic to flow, AS_B has to accept this routing information and use it. It is AS_B's privilege to either use or disregard the information that it receives from AS_A about NET1's reachability. AS_B might decide not to use this information, in cases it does not wish to send traffic to NET1 or if it considers more appropriate to use another route to reach NET1.

Inter-AS routing policies are often complex and motivated by the need to balance the traffic on links with other ASs and to reduce the cost of carrying traffic on these links. These requirements rely on the connectivity of an AS with other ASs and its AS relationships. Typically, an AS aims to optimize the way traffic enters or leaves its network based on its business interests. For example, content-providers will try to optimize the way traffic leaves their networks. On the other hand, access-providers that serve small and medium enterprises, dialup, or other connection services, usually wish to optimize how Internet traffic enters their networks. Finally, a transit AS, AS that connects with neighboring peer ASs, will try to balance the traffic on the links with its peers.

Ideally, the announcement and acceptance policies of AS_A and AS_B are symmetrical. The BGP protocol allows each AS to choose its own administrative policy for selecting routes and propagating reachability information to other routers. However, reaching a destination in another AS requires the use of resources or routers along the route. Hence, routing policies are constrained by administrative policies and by the commercial agreements between ASs. For instance, an AS would often set its policy so that it does not export routes of its rivals or it would prefer to use the resources of a particular AS over other ASs.

AS relationships can be classified as customer-provider and peering. In a customer-provider relationship, the customer AS typically belongs to a smaller administrative domain that pays a larger administrative domain where the provider AS belongs to, for access to the rest of the Internet. In a peer relationship, the two peering ASs typically belong to administrative domains of comparable size and they exchange traffic between their respective customers. Via route advertisements an AS sets its export policies, according to its relationships with the neighboring ASs. AS relationships may be translated into the following rules that govern route advertisements [30]:

- Exporting to a customer: An AS can export routes originated by it, routes originated by its customers, and routes learned from other providers or peers.
- Exporting to a provider: In exchanging routing information with a provider, an AS can export routes originated by it and tag the “no-export” BGP communities attribute to the routes that do not need to be propagated beyond the provider’s AS. The BGP communities attribute is an optional attribute that can be attached to routes. An AS can also export to a provider routes originated by its customers.
- Exporting to a peer: an AS can export routes originated by it and tag the “no-export” BGP communities attribute to the routes that do not need to be propagated beyond

the peer's AS. An AS can also export to a peer routes learned from other providers or peers.

ASs can have rather complicated relationships. For instance, two ASs operated by the same institution may have a customer-provider relationship where each AS offers transit service [18]. Some AS pairs may have back up relationships to provide connectivity in case of failure [19]. Other AS pairs may peer indirectly through a transit AS [20]. Sometimes, an AS pair may have different relationship for certain block of IP address. For example, an AS in the United States may be a customer of an AS in Europe for some destinations and a peer for others. Router misconfiguration may cause a violation of the export rules that results in erroneous AS relationships. For instance, a peer may mistakenly export advertisements learned from one provider to another. These exceptions are usually rare and we assume that only a small fraction of the AS pairs do not have the traditional provider-customer and peer-peer relationships [40].

Routes transmitting back and forth between source and destination ASs in the Internet may be different or asymmetrical due to control of incoming traffic. Recently, F. Wang et al., [38] noticed that ASs may announce their prefixes to only a subset of their providers. If a provider receives a prefix propagated by a customer via a peer path, instead of a customer path, the prefix is called "selective announced prefix"(SA prefix). The selective announcement routing policies employed by customer B can be observed at provider A as shown in Figure 2.3. Customer B announces prefix p to provider C but not to provider A. In the BGP table of provider A, prefix p is received from its peer C. The primary reason for the SA prefix is to control the incoming traffic. These routing policies imply that there are less available paths in the Internet than shown in the AS connectivity graph. Customers can optimize their inbound traffic by applying selective announcement policies,. Their inbound and outbound traffic might be asymmetric. Providers may find

that traffic between their customers has to be forwarded to the rest of the Internet via heavy loaded peer links.




Figure 2.3 An example of a selected announcement prefix.

2.6 Geographic properties of the Internet topology

Attempts to model the Internet structure have often made explicit or implicit assumptions about the generated network's geometry. For example, in Waxman's model [39], two assumptions are made: (1) network nodes are placed uniformly at random in the plane; and (2) the probability that two nodes are directly connected is an exponentially declining function of the separation distance. Other models have implicitly assumed no important underlying geographic properties of the network and that the patterns of connectivity are only influenced by topological factors [7], [13], [26].

Lakhina et al., [23] observed that geographic location of the Internet resources plays an important role in shaping the Internet topology factors, after studying the Internet's physical structure concerning the geographical location of its components: routers, links and ASs. For example, they found in the Waxman assumptions, that assumption 1 uniform distribution of nodes is inaccurate because the distribution pattern of nodes is highly irregular. However, assumption 2 is valid because the connectivity patterns of nodes show a strong relationship to distance between nodes. Furthermore,

they predicted that the next generation of topology generators would be geographically based. To provide guidelines to the development of these geographically-driven generation methods, Lakhina et al., analyzed a dataset from CAIDA and a dataset collected by *traceroute* in the Scan Project [36]. They showed that the connection patterns between routers are strongly related to geographical distance. Also, the result shows that the number of distinct locations spanned by an AS is strongly correlated with the number of interfaces (routers) and degree in the AS graph.

AS geography is also considered as a controlling factor when constructing the modified HOT model for Internet growth [10]. In [10], the authors argued that large ASs are more likely to acquire new ASs because of their proximity to the topological core of the AS graph and the geographic diversity of their Point of Presence (POP) infrastructure. POP is an access point to the Internet. The number of POPs that an ISP has is sometimes used as a measure of the growth rate of the ISP's size. This acquisition of new ASs in turn enables large ASs to build up their POP infrastructure more aggressively than small ASs. Hence, it is plausible that high variability in the number of POPs per AS may cause AS degrees to exhibit high variability.

Geography has been used to analyze various aspects of the Internet routing. Subramanian et al., [33] found that the circuitousness (how circuitous a route is) of routes in the Internet depends on the geographic and network locations of the end-hosts. Circuitousness of a path is strongly correlated to minimum delay characteristics, which is an important network performance metric.

Chapter 3 Spectral Graph Theory

In this Chapter, we present basics of spectral graph theory and eigen-analysis of the Internet topology.

3.1 Spectra of a graph

A graph $G(V, E)$ is a set of vertices V connected by a set of edges E . A *loop* is an edge with both of its vertices identical. If multiple edges are allowed, the graph is called *multigraphs*. Multigraphs may contain loops. An example of multigraph is illustrated in Figure 3.1.




Figure 3.1 An example of a multigraph.

An Internet AS graph considered in the thesis represents a set of ASs connected via logical links. The number of edges incident to a node in an undirected graph is called the *degree* of the node. A digraph is a directed graph with a set of nodes connected by a set of directed links. In digraphs, *indegree* and *outdegree* of a node indicate the number of links that are directed to and out of a node, respectively. Two nodes are called *adjacent* if they are connected by a link. We call a graph *complete graph* if any pair of

graph vertices is connected with an edge in the graph. A complete graph with n vertices is denoted K_n . For instance, K_2 , K_3 and K_6 are shown in Figure 3.2.




Figure 3.2 An example of a complete graph.

A *path* in a graph is a sequence $\{x_1, x_2, \dots, x_n\}$ such that $(x_1, x_2), (x_2, x_3), \dots, (x_{n-1}, x_n)$ are graph edges. A multigraph is *connected* if any two of its vertices are connected by a path. A multigraph is *disconnected* if it is not connected, and hence, it consists of two or more parts called *connected components*. In a disconnected multigraph, two vertices are in different connected components if they cannot be joined by a path.

A *bipartite* graph is a set of graph vertices decomposed into two disjoint sets such that no two vertices within one set are adjacent. An example is shown in Figure 3.3.




Figure 3.3 An example of a bipartite graph.

A graph $H (V', E')$ is said to be a *subgraph* of the graph $G (V, E)$ if $V' \subset V$ and $E' \subset E$. If all the edges that connect all the vertices in V' are included in E' , H is called an *induced subgraph*. If an induced subgraph is itself a complete graph, it is called *clique*. For example, the graph shown in Figure 3.4 (b) is an induced subgraph of the graph K_6 shown in Figure 3.4 (a). It is also a clique.




Figure 3.4 An example of a clique graph.

A graph G can be represented by its adjacency matrix $A(G)$:

$$A(G)_{i,j} = \begin{cases} 1 & \text{if } i \text{ and } j \text{ are connected,} \\ 0 & \text{otherwise.} \end{cases} \quad (3.1)$$

A diagonal matrix $D(G)$ associated with $A(G)$ is a matrix with row-sums of $A(G)$ along the diagonal and the rest of the vertices zero.

$$D_{ij}(G) = \begin{cases} \sum_j A(i, j) & \text{if } i = j, \\ 0 & \text{otherwise,} \end{cases} \quad (3.2)$$

$D(G)$ indicates the connectivity degree of each node. The Laplacian matrix is defined as $L(G) = D(G) - A(G)$:

$$L_{i,j}(G) = \begin{cases} D_{ij}(G) & \text{if } i = j, \\ -1 & \text{if } i \text{ and } j \text{ are adjacent,} \\ 0 & \text{otherwise,} \end{cases} \quad (3.3)$$

Eigenvalues of a matrix M are defined as numbers λ satisfying $Mx = \lambda x$ for a non-zero vector x . Vector x is called an eigenvector of the matrix M associated with eigenvalue λ . The collection of all eigenvalues is called a *spectrum*. For example, the adjacency matrix A , the diagonal D and the Laplacian matrix L of the graph shown in Figure 3.5 are:

$$A = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix} \quad D = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad L = \begin{pmatrix} 1 & -1 & 0 & 0 \\ -1 & 2 & -1 & 0 \\ 0 & -1 & 2 & -1 \\ 0 & 0 & -1 & 1 \end{pmatrix}. \quad (3.4)$$

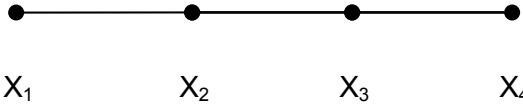


Figure 3.5 Adjacency matrix A , Diagonal matrix D , and Laplacian matrix L of a sample graph.

For the graph shown in Figure 3.5, the spectrum s ($\lambda_1 \dots \lambda_4$) and associated eigenvectors ($y_1 \dots y_4$) are given in Table 3.1:

Table 3.1 Spectrum and eigenvectors of a sample graph.

	Element 1	Element 2	Element 3	Element 4
s	-1.6180	-0.6180	0.6180	1.6180
y_1	0.3717	-0.6015	-0.6015	0.3717
y_2	-0.6015	0.3717	-0.3717	0.6015
y_3	0.6015	0.3717	0.3717	0.6015
y_4	-0.3717	-0.6015	0.6015	0.3717

The second smallest eigenvalue (a signed value) of L is called the *algebraic connectivity* of a graph [17]. The eigenvalue conveys many properties related to

connectivity of a graph. The vector of the eigenvalue, the second eigenvector, is called the vector of the algebraic connectivity.

Assume $G = (V, E)$ is a connected graph and y the second eigenvector. The elements of this eigenvector are assigned to the vertices of G so that each vertex of G has a corresponding eigenvalue in the eigenvector. The assignment can be considered as valuations of the vertices of G . Fiedler called the process “characteristic valuation” of G and proved that characteristic valuation can be useful in partition problems in graph theory.

Assume a vertex k of G is common to more than one subgraph. A subgraph each contains a subset of edges and all vertices adjacent to them. All these subgraphs together will cover all the edges of G . Let G_0, G_1, \dots, G_r be all components of the graph obtained from G by removing vertex k and all adjacent edges. Then,

- If the k th element in eigenvector y $y_k > 0$, then exactly one of the components of G contains a vertex with negative value in y . For all vertices s in the remaining components, $y_s > y_k$.
- If $y_k = 0$ and there is a component of G containing vertices with both positive and negative value, then there is exactly one such component and all remaining components only have vertices with value zero.
- If $y_k = 0$ and no component contains both positively and negatively valued vertices, then each component of G contains either only positively valued, or negatively valued or only zero valued vertices.

Laplacian matrix is often normalized. The normalized Laplacian matrix $N(G)$ is defined as [12]:

$$N_{i,j}(G) = \begin{cases} 1 & \text{if } i = j \text{ and } d_i \neq 0 \\ -\frac{1}{\sqrt{d_i d_j}} & \text{if } i \text{ and } j \text{ are adjacent} \\ 0 & \text{otherwise,} \end{cases} \quad (3.5)$$

where d_i and d_j are the degrees of node i and j , respectively.

The advantage of definition (3.5) is due to the fact that this normalized Laplacian is consistent with the eigenvalues both in spectral geometry and in stochastic processes. Many results that were only related to regular graphs can be generalized to all graphs. An r -regular graph is a graph with all vertex degrees equal to r . Therefore, the normalized Laplacian provides a “coherent treatment” for a general graph [12].

We denote eigenvalues of normalized Laplacian matrix N by $0 = \lambda_0 \leq \lambda_1 \leq \dots \leq \lambda_{n-1}$.

1. Basic facts about N are [12]:

- $0 \leq \lambda_i \leq 2$, for all $i \leq n-1$, with $\lambda_{n-1} = 2$ if and only if a connected component of G is bipartite and nontrivial.
- $\sum_i \lambda_i \leq n$, equality holds if and only if G has no isolated vertices.
- For $n \geq 2$, $\lambda_1 \leq \frac{n}{n-1}$. The equality holds if and only if G is the complete graph on n vertices. Furthermore, for a graph G without isolated vertices,

$$\lambda_{n-1} \geq \frac{n}{n-1}. \quad (3.6)$$

- $\lambda_1 \leq 1$ if a graph is not a complete graph.
- If G is connected, then $\lambda_1 > 0$. If $\lambda_{i+1} \neq 0$, then G has exactly $i+1$ connected components.

The expansion property of graphs means that each subset S of vertices must have “many” neighbors. That is, the neighborhood set $N(S) = \{y \mid y \text{ is adjacent to some } x \in S\}$ is “large” compared to the size of S . Quantitative description of “many” or “large” can be found in [12]. An expander graphs is defined as a regular graph G of n vertices is a c -expander if every subset S of $V(G)$ satisfies:

$$|N(S)| \geq \left(1 - \frac{|S|}{n}\right) |S|, \quad (3.7)$$

where the constant θ is called the expander coefficient.

An expander graph allows us to locate clusters (any “small” subset of vertices that has a relatively “large” neighborhood in a graph). In communication networks, a cluster may indicate better connections and higher routing abilities for vertices within the cluster.

For a subset S of $V(G)$, the vertices of G , we define volume of S , $vol S$, to be the sum of degrees of the vertices in S [12]:

$$vol S = \sum_{x \in S} d_x, \quad \text{for } S \subseteq V(G). \quad (3.8)$$

It can be seen that $vol S$ also indicates the number of edges in S .

The key to the success of constructing expander graphs is their relationship with eigenvalues. We are interested in a lower bound of expander graphs [12]:

Suppose that G is not a complete graph. For $S \subseteq V(G)$, the neighborhood $N(S)$ satisfies

$$\begin{aligned}
\frac{\text{vol } N(S)}{\text{vol } S} &> \frac{1}{\bar{\lambda}^2 + (1 - \bar{\lambda}^2) \frac{\text{vol } S}{\text{vol } G}} \\
&= \frac{1}{1 - (1 - \bar{\lambda}^2) \frac{\text{vol } \bar{S}}{\text{vol } G}}, \tag{3.9}
\end{aligned}$$

where $\bar{\lambda} = \max_{i \neq 0} |1 - \lambda_i|$. The lower bound suggests a strong connection between the largest eigenvalues of N and clusters in G .

3.2 Eigen-analysis

Eigenvalues associated with a network graph are closely related to important topological features, such as diameter of the network, presence of cohesive clusters, long paths and bottlenecks, and the randomness of the network graph. Faloutsos et al., [15] have considered the first 20 largest eigenvalues of the adjacency matrix of a graph.

Vukadinović et al., [37] reported that the normalized Laplacian spectrum (n/s) of the Internet topology on AS level (AS graph) is invariant regardless of the exponential growth of the Internet. They found remarkably similar plots of the n/s for real Internet AS-level data spanning several years. The same consistency was also reported for synthetically generated graphs with various numbers of nodes. Hence, n/s can distinguish between AS graphs and synthetically generated graphs and is an excellent candidate as a fingerprint of Internet graphs.

Mihail et al., [26] used the eigenvectors corresponding to the largest eigenvalues of the Laplacian matrix to find clusters of ASs with certain characteristics, such as geographic locations or business interests. They used a spectral filtering method that separates clusters by examining the eigenvectors. The spectral filtering method for a symmetric matrix A are:

- Compute several largest eigenvalues of A and the corresponding eigenvectors.
- Sort nodes according to the elements in the eigenvectors. Nodes in the graph directly relate to the index of elements in the eigenvectors of the graph's matrix. In [26], the matrix is a transformed, normalized adjacent matrix of the graph.
- Cut towards the most positive end or towards the most negative end of the sorted elements, with special preference to sharp jumps. These selected groups are candidates for clustering.

Chapter 4 Analysis of Internet Data

4.1 Observations of Internet topology datasets

We use datasets collected on a typical day in May, 2003 from Route Views and RIPE. A typical day implies a day that no global-scale virus attacks have occurred. It is observed that Internet topology data are greatly affected (size of data shrunk dramatically) when virus attacks occurred. Since the dynamics and the evolution of Internet topology are not our analysis issues, we do not use data collected over a period of time.

During data preprocessing, we extracted AS routes from data files and segmented the AS_PATH attribute of the routes into AS pairs. AS pair consists of two ASs adjacent to each other in an AS_PATH. We remove pair duplicates and consider the direction of AS pairs. For example, in two AS paths 12-222-45-34 and 222-12-45-34, we can collect 5 AS pairs: 12-222, 222-45, 45-34, 222-12 and 12-45. We consider the direction of a route to reflect the fact that AS connectivity does not guarantee AS reachability.

Table 4.1 Statistics of Route Views and RIPE datasets.

	Route Views	RIPE
Number of AS routes	6,398,912	6,375,028
Number of probed ASs	15,418	15,433
Number of AS pairs	34,878	35,225

The number of the AS routes and AS pairs is shown in Table 4.1. The Route Views dataset consists of 15,418 assigned ASs, and RIPE has 15,433. The two datasets contain similar information. Of the collected ASs in each dataset, 15,369 matching ASs are found. The two datasets probed almost the same set of ASs in the Internet with only

0.3% differences. 29,477 AS pairs can be found in both datasets. This represents approximated 85% of the AS pairs in the Route Views dataset and 84% in the RIPE dataset.

We ordered the ASs according to connectivity level (degree of an AS node). The result is shown in Table 4.2. Fourteen of twenty ASs with the largest node degrees in both datasets are identical, i.e., 70% of the core ASs (ASs with the largest degrees) are identical in the two datasets. Furthermore, core ASs in the Route Views dataset have larger degrees than core ASs in RIPE.

Table 4.2 Assigned numbers of twenty ASs with the largest node degrees.

Rank of degree	Route Views		RIPE	
	ASN	Degree	ASN	Degree
1	701	2595	701	2448
2	1239	2569	1239	1784
3	7018	1999	7018	1638
4	3561	1036	209	861
5	1	999	3561	705
6	209	863	3356	673
7	3356	662	3549	612
8	3549	617	702	580
9	702	562	2914	561
10	2914	556	1	489
11	6461	498	4589	482
12	4513	468	6461	476
13	4323	315	8220	450
14	16631	294	3303	429
15	6347	291	13237	412
16	8220	289	6730	313
17	3257	277	4323	305
18	4766	263	3257	305
19	3786	263	16631	296
20	7132	258	6347	281

4.2 Spectral analysis of the AS Internet topology

We analyze the two datasets spectrally. We considered ASs with the first 30,000 assigned AS numbers. Operations on the matrix from the Internet AS graph, which can contain more than 65,000 ASs (vertices), will be overwhelming. As shown in Figure 4.1,

in both datasets, most active ASs, ASs that have degree larger than 0, occupy the first 30,000 AS numbers. Hence, we can only consider the first 30,000 ASs in order to minimize the computation without jeopardizing the accuracy of the result.




Figure 4.1 AS degree distribution in Route Views and RIPE datasets.

We perform the same characteristic valuation process to analyze the 30,000 ASs and depict the result in Figure 4.3. This figure is built to easily show the status of connectivity and clustering of ASs. We use a simple sample to illustrate how the characteristic valuation process works with the second smallest eigenvector and how to interpret the resulting figure. Consider a graph with 4 nodes: AS1, AS2, AS3, and AS4. The second smallest eigenvector calculated in the graph is $[0.1, 0.3, -0.2, 0]$. We assign elements of the eigenvector to the nodes in the order of the index and the result is: $[AS1(0.1), AS2(0.3), AS3(-0.2), AS4(0)]$. The ASs are then sorted by their element value in an ascending order. The resulting AS vector is $[AS3, AS4, AS1, AS2]$. Assume that only nodes AS3 and AS1 are connected (with degrees larger than 0) in the graph. The resulting figure is shown in Figure 4.2. The X axis in Figure 4.2 refers to the index of the element (an AS) in the final AS vector. The Y axis indicates the connectivity status of

corresponding ASs. If the AS is connected, its value of the connectivity status is 1. It is 0 if the AS is isolated. Also, ASs have close element values will stay closer in the figure due to the sorting.






Figure 4.2 An example demonstrates how characteristics valuation process works.

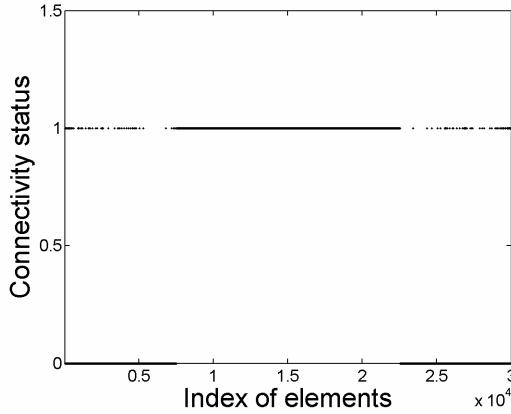
Figure 4.3 shows the connectivity status of elements of the second smallest and the largest eigenvectors in Route Views and in RIPE respectively.




(a) 2nd smallest eigenvector in Route Views



(b) largest eigenvector in Route Views



(c) 2nd smallest eigenvector in RIPE



(d) largest eigenvector in RIPE

Figure 4.3 Spectral views of AS connectivity in two datasets.

The result shows that eigenvectors corresponding to the second smallest eigenvalues tend to partition data into connected and non-connected ASs. Among the 30,000 ASs, ~15,000 ASs with a degree larger than zero (connected ASs) are grouped together as shown in Figure 4.3 (a) and (c). The cores of the connected ASs, the center parts in the figures, of the two datasets are similar.

Eigenvectors corresponding to the largest eigenvalue reveal highly different clustering characteristics in two datasets, even though the two datasets collected almost the same set of ASs and most connections (AS pairs) among these ASs are identical. In

Figure 4.3 (b), two large clusters are visible in the Route Views dataset. In Figure 4.3 (d), values of 1 and 0 are tightly interwoven because the RIPE dataset has a larger number of smaller clusters compared to Route Views. (Because of the large sample size of 30,000 points in Figure 4.3 (d), the seemingly connected lines are actually composed of many disconnected segments). In the RIPE dataset, nodes are relatively dispersed and many small clusters are scattering throughout the space.

We choose a small cluster of ASs in RIPE in Figure 4.3 (d) and search for the locations of the selected ASs in Route Views in Table 4.3 (c). ASs are clustered together if they have close element values after the characteristic valuation process. The result shows that ASs in the cluster (shown in Table 4.3) are separated into two large clusters in Route Views (shown in Table 4.4). The reasons for the distinct clustering characteristics in the two datasets may be, as shown in Table 4.2, core ASs in Route Views have larger degrees than core ASs in RIPE. Therefore, core ASs in Route Views tend to connect a larger number of smaller ASs, while ASs in RIPE dataset are more likely to be dispersed.

Table 4.3 An example of a small cluster in the RIPE dataset.

Element value	ASN	Location
9.87E-02	21032	Germany (TELTA Citynetz Eberswalde)
1.02E-01	2450	France (INRIA-Rocquencourt)
1.02E-01	2426	France (RUBIS Metropolitan Area Network)
1.09E-01	9651	Australia (DOT Communications)
1.09E-01	9652	Australia (ECN Internet)
1.10E-01	16906	Latin America (El Salvador Network)
1.18E-01	13136	Netherlands (Interstroom Informatietechnologie)
1.20E-01	25125	Israel (Israel Local Authorities Data Processing Center)
1.24E-01	21922	USA (Webnet Memphis, Inc.)
1.27E-01	14708	Latin America (WebHost)
1.29E-01	24807	UK (Infocom UK Ltd)
1.61E-01	7566	Australia (Teragen Internet Solutions)
1.77E-01	20908	Poland (CR-MEDIA)
2.01E-01	14647	USA (Network O.S., Inc.)

Table 4.4 ASs (belong to one cluster in RIPE dataset) in Route Views dataset are separated into two clusters by element values.

Element value	ASN		Element value	ASN
-1.06E-05	21032		5.35E-05	2450
-1.38E-05	9651		5.35E-05	2426
-1.61E-05	25125		4.77E-07	9652
-7.58E-05	24807		1.46E-05	16906
-2.11E-05	7566		8.13E-06	13136
			3.1E-04	21922
			3.17E-05	14708
			2.51E-06	20908
			6.56E-05	14647

Network clustering is a major performance metric and has been extensively studied [26]. Our results show that the two datasets may exhibit different clustering characteristics of the Internet. Since the two datasets are widely used, future Internet topology study specifically regarding clusters may need to pay attention to the dataset being used.

4.3 Reverse pairs

The geographical difference between participating ASs in Route Views and RIPE may reveal the inter-AS routing policies employed by network operators. Participating ASs in the two datasets differ geographically. As can be seen from the list of participating ASs in the Route Views project (Appendix A-1), more than 80% ASs are located in North America. Over 90% ASs found in the RIPE dataset (Appendix A-2) reside in Europe.

Most participating ASs in both datasets belong to access-providers. These access-providers may prefer that incoming traffic be localized to their specific geographic areas. Access-providers usually serve local small and medium enterprise, dialup and other connection services. They usually wish to optimize the way Internet traffic enters their networks.

Routing policies on incoming traffic of ASs will influence AS connectivity. For example, assume for a pair of connected ASs A and B, A wishes to control incoming traffic from B for certain reasons. A can accomplish this goal by not sending B its route advertisement. Nevertheless, B may still send its route advertisement to A. Consequently a unidirectional route from A to B is formed in routing tables. Hence, unidirectional routes may exist due to the routing control of the incoming traffic.

We are interested in certain unidirectional routes from North America to Europe or from Europe to North America. The Route Views and RIPE datasets were collected from geographically different locations. Most of the participating ASs in Route Views are located in North America while most of the ASs in RIPE reside in Europe. The participating ASs in Route Views are access-providers that tend to limit incoming traffic from Europe, because these North America ASs typically serve customers in North America. Consequently, the originating ASs tend to select ASs in North America as their next hops of routes in routing tables. In turn, the same process may very likely be performed by the selected ASs until ASs in North America cannot be found, then ASs in other places including Europe are selected. Hence, in Route Views, more unidirectional routes originating from North America to Europe are expected to be found. For dataset in RIPE, more unidirectional routes originating from Europe to North America are expected to be located.

We introduce a new metric called the reverse pair to facilitate the analysis of the routing design practice on incoming traffic in the Route Views and RIPE datasets. We call two ASs, A and B, a *reverse pair* $R(A, B)$ in two dataset S and T if A and B satisfy:

$$\{(A, B) \mid (A-B) \in (\text{AS pairs in } S) \text{ and } (A-B) \notin (\text{AS pairs in } T) \\ \text{and } (B-A) \in (\text{AS pairs in } T) \text{ and } (B-A) \notin (\text{AS pairs in } S)\}. \quad (4.1)$$

Where S (or T) is the Route Views dataset and T (or S) is the RIPE dataset. Note that AS pair $A-B$ is different from AS pair $B-A$. We conjecture that reverse pairs may indicate the special unidirectional routes that could suggest geographically related routing policies on incoming traffic. To test this hypothesis, we show that reverse pairs in dataset of Route Views have more ASs originating from North America and that reverse pairs in the RIPE dataset have more ASs originating from Europe.

We found 558 reverse pairs in the Route Views and RIPE datasets that we analyzed. They represent approximately 1.60% and 1.58% of all the AS pairs in Route Views and RIPE, respectively. There are 189 AS in the 558 reverse pairs. “Degrees” of these ASs range from 1 to 38. Note that the “degrees” here are calculated only by counting links among ASs belong to reverse pairs. Both datasets have approximately 85% of AS pairs in common. This implies that the remaining 15% AS pairs are distinct. Hence, the reverse pairs proportion in the distinct AS pairs is not negligible.

We consider the “outdegrees” of ASs that belong to reverse pairs in order to infer originating ASs in two datasets. For example, an AS that is the originating ASs of two reverse pairs will have an “outdegree” equal to two. In the analysis, we calculate the “outdegrees” and “indegrees” of ASs of reverse pairs in RIPE and Route Views. The results of core ASs (ASs with “degree” total larger or equal to 10) are given in Tables 4.5 (a) and (b) respectively.

Table 4.5 ASs with degree total larger or equal to 10 among the reverse pairs in (a) RIPE, (b) Route Views. In the column of Location, EU refers to the AS is in Europe, NA refers to in North America, and ASIA is in Asia.

(a)				(b)			
ASN	Out-degree	In-degree	Location	ASN	Out-degree	In-degree	Location
3303	35	3	EU	3257	29	1	EU
6730	27	3	EU	6461	26	0	NA
3320	24	3	EU	4513	24	0	NA
4589	21	1	EU	3356	22	0	NA
15412	20	1	EU	3561	18	0	NA
3300	19	1	EU	12956	17	0	EU
4200	18	1	NA	3246	16	0	EU
5400	18	3	EU	3549	15	0	NA
8220	17	2	EU	4637	15	0	ASIA
13237	16	2	EU	1239	14	0	NA
297	15	0	NA	8001	14	0	NA
6762	15	3	EU	2516	13	0	ASIA
13129	14	0	EU	2497	12	0	NA
2529	13	1	EU	2914	12	0	NA
286	12	1	EU	7911	12	0	NA
1759	10	1	EU	3333	11	0	EU
6467	10	1	EU	702	10	8	NA
				1299	10	3	EU
				5511	10	0	EU
				6453	10	0	NA

Table 4.5 shows that reverse pairs in Route Views dataset have more ASs originating from North America while reverse pairs in RIPE dataset have more ASs from Europe. In Table 4.5 (a), the “Outdegree” column indicates the number of occurrences the AS is an originating AS among reverse pairs in RIPE. The “Outdegree” column in Table 4.5 (b) shows the number of occurrences the AS is an originating AS in Route Views. We observe that most originating ASs (15 out of 17, approx. 88%) in the RIPE dataset are located in Europe. The majority of originating ASs (12 out of 20, 60%) in the Route Views dataset are in North America. Hence, reverse pairs may be used to indicate the unidirectional routes that suggest geographically based routing policies on the incoming traffic. Most core ASs belong to reverse pairs are large ASs (with degrees

larger or much larger than 100, as shown in Table 4.2). This may be attributed to that large ASs often have regional routing policies [22].

We also construct routes consisting of reverse pairs. We connected two reverse pairs if the originating AS in one pair is the ending AS of the other pair. Isolated routes are isolated reverse pairs. The results in Table 4.6 show that routes built in this way do not often have more than 1 hop and never exceed 3 hops.

Table 4.6 Statistics of routes built from reverse pairs.

Number of Reverse pairs	Total number of routes	Number of isolated routes	Number of routes with more than 1 hops
558	503	11	56

Routes with large number of hops built from reverse pairs are usually geographically dispersed probably because reverse pairs exist in international links. All routes with 3 hops are shown in Table 4.7. An example of the route shown in the first row of Table 4.7 illustrates how dispersed this type of route may be as shown in Figure 4.4.

Table 4.7 Routes built from reverse pairs with the maximum number of hops 3.

1103 (EU)	702 (NA)	6762 (EU)	1239 (NA)
8406 (EU)	8210 (EU)	4200 (NA)	3549 (NA)
5417 (EU)	702 (NA)	1299 (EU)	8297 (EU)
6893 (EU)	12541 (EU)	1273 (EU)	4513 (NA)
12381 (EU)	1653 (EU)	2603 (EU)	3257 (EU)
28764 (EU)	24745 (EU)	12713 (EU)	3561 (NA)
15623 (EU)	12755 (EU)	8220 (EU)	3356 (NA)




Figure 4.4 An example of reverse pair path.

As the data collecting techniques get more advanced and the covered Internet areas more complete, the two large Internet datasets from Route Views and RIPE will contain additional valuable information about the Internet topology. Our analysis suggests that data from the Route Views and RIPE datasets are collected from geographically different ASs, and hence, the two datasets may assist in analyzing geographically related routing policies on incoming traffic. Our results also show that the effect of these routing policies is not negligible in Internet network operators.

Chapter 5 Conclusions

Internet topology is important to the Internet research and development. Due to the underlying complex mechanisms and the extremely large size of the Internet, studies of Internet topology rely either on limited Internet AS datasets or on employing synthetic topology generators. The results of the studies may be affected by the characteristics of the data sources being used.

In this thesis, we analyzed two major available Internet topology datasets from Route Views and RIPE to address concerns related to the two datasets. We analyzed the distinct clustering features of the two datasets spectrally. After analyzing the eigenvectors of the largest eigenvalues from the two datasets, we observed that the two datasets exhibit distinct clustering characteristics. Network clustering is a major performance metric [26]. A cluster in the network indicates better connections and routing abilities of the nodes within the cluster. Our result suggests that Internet topology study regarding Internet clustering need to pay attention to the dataset being used. For example, it is likely that network clustering studies rely on data from Route Views will consider the Internet presently has better connections and routing abilities. The reason to the conclusion may be that the Internet graph (shown in Figure 4.3 (b)) from the Route Views dataset consists of two giant clusters in our spectral analysis. However, studies using RIPE dataset may reach different conclusion.

We also propose the notion of “reverse pairs” and use it as a new metric to analyze the datasets in order to study Internet routing policies. Our analysis suggests that locating “reverse pairs” in the Route Views and RIPE datasets may assist in

analyzing routing policies on incoming traffic in the Internet because they are collected from geographically different locations.

Appendix A List of participating ASs in the Route Views and RIPE datasets

Appendix A-1 Participating ASs in Route Views

Name of AS organization	Location	IP of AS router	AS assigned number
Abilene		198.32.8.252	11537
Abilene	Indiana	198.32.8.196	11537
Accretive	PAO	207.246.129.6	11608
Accretive	SEA	207.246.129.14	11608
AOL	NoVa	66.185.128.48	1668
APAN/tppr-tokyo	ORD	203.181.248.233	7660
Army Research Laboratory		192.12.65.1	13
ATT	CA	12.0.1.63	7018
ATT/Canada	east	216.191.65.118	15290
ATT/Canada	west	216.191.65.126	15290
Blackrose.org	Ann Arbor	204.212.44.131	234
Broadwing	ADDS	216.140.14.186	6395
Broadwing	MaeEast	216.140.8.59	6395
Broadwing	MaeWest	216.140.2.59	6395
C&W	Santa Clara	208.172.146.2	3561
CA*net3		205.189.32.153	6509
Carrier1	AMS	212.4.193.253	8918
COMindico	AU	203.194.0.12	9942
Digex	VA	209.19.207.70	2548
ELI	MAE-EAST	208.186.154.35	5650
ELI	MAE-WEST	208.186.154.36	5650
Epoch	PAO	155.229.0.36	4565
ESnet	GA	134.55.20.229	293
France Telecom Backbone	NYC	193.251.128.22	5511
Global Crossing	SEA	208.51.113.253	3549
GLOBIX	LINX	195.66.224.82	4513
GLOBIX	New York	209.10.12.28	4513
GLOBIX	ORD	209.10.12.125	4513
GLOBIX	PAO	209.10.12.156	4513
GT Group Telecom Service	ORD	216.18.63.137	6539
Hurricane Electric	DCA	216.218.252.152	6939
Hurricane Electric	PAO	216.218.252.145	6939
IJJ	Japan	202.232.0.2	2497
IP-PLUS	ZRH	164.128.32.11	3303

Name of AS organization	Location	IP of AS router	AS assigned number
Jippii	ESPANIX	62.164.11.10	8782
KPNE	AMSIX	134.222.85.45	286
Level3	DEN	209.244.2.115	3356
LINX	London	195.66.232.239	5459
LINX	London	195.66.232.254	5459
M-Root	Japan	202.249.2.86	7500
MFN	PAO	209.249.254.19	6461
MFS/MAE-Lab	SJC	204.29.239.1	6066
MSN	PAO	207.46.32.32	8075
NCSA	MDW	141.142.12.1	1224
Net Access	NYC	209.123.12.51	8001
Nether.net	Ann Arbor	204.42.253.253	267
netINS	DSM	167.142.3.6	5056
Netrail	ORD	205.215.45.50	4006
Port80	STO	217.75.96.60	16150
RCN	DCA	207.172.6.162	6079
RCN	PAO	207.172.6.227	6079
RIPE NCC	Amsterdam	193.0.0.56	3333
RUSnet	MOW	194.85.4.249	3277
Sprint	Stockton	144.228.241.81	1239
Sprint/Canada	YYZ	206.186.255.223	2493
STARTAP		206.220.240.95	10764
TDC	NYC	195.249.0.135	3292
TDS Telecom	MSN	64.50.230.1	4181
TDS Telecom	MSN	64.50.230.2	4181
telefonica	New York/GRTNYCCC2	213.140.32.146	12956
Teleglobe	London	195.219.96.239	8297
Teleglobe	PAIX	207.45.223.244	6453
Telia	NYC	213.248.83.240	1299
Telstra	Sydney	203.62.252.26	1221
Telus	Calgary	154.11.98.18	852
Telus	Tornoto	154.11.63.86	852
The University of Waikato	AKL	130.217.2.25	681
Tiscali	PAR	213.200.87.254	3257
TouchAmerica	PDX	157.130.182.254	19092
UONet	Oregon	198.32.162.1	3582
UUNET	Africa	196.7.106.245	2905
Verio	CA	129.250.0.11	2914
Verio	VA	129.250.0.85	2914
WCICABLE	Hillsboro OR	209.161.175.4	14608
Williams	PAO	64.200.199.4	7911
Williams	SFO	64.200.199.3	7911
XO	SJC	65.106.7.139	2828

Appendix A-2 RRC information in RIPE

Prefix	RRC	Location
195.80.224.0/24	RRC00 - RIPE NCC	Amsterdam, NL
195.80.225.0/24	RRC01 – LINX	London, UK
195.80.226.0/24	RRC02 – SFINX	Paris, FR
195.80.227.0/24	RRC03 - AMS-IX	Amsterdam, NL
195.80.228.0/24	RRC04 – CIXP	Geneva, CH
195.80.229.0/24	RRC05 – VIX	Vienna, AT
195.80.230.0/24	RRC06 – NSPIX2	Otematchi, JP
195.80.231.0/24	RRC07 – Netnod-IX	Stockholm, SE
195.80.232.0/24	RRC08 - MAE-WEST	San Jose, CA, US

RRC00 - RIPE NCC Peer List:

AS513	CERN - European Organization for Nuclear Research		192.65.184.3
AS1103	SURFnet	The Netherlands	195.69.144.34
AS2858	EUnet Test AS		194.109.197.245
AS2914	No description available		129.250.0.232
AS3257	Tiscali Intl Network		195.69.144.85
AS3333	RIPE NCC		193.0.0.56
AS3549	Global Crossing Ltd.		195.66.224.112
AS3549	Global Crossing Ltd.		64.211.147.146
AS4608	Asia Pacific Network Information Center Pty. Ltd.		202.12.29.64
AS4777	Asia Pacific Network Information Centre		202.12.28.190
AS7018	No description available		12.0.1.63
AS9177	SOLPA AG		212.47.190.1
AS13129	Global Access Telecommunications Inc.		212.20.151.234

RRC01 Peer List:

AS286	KPN Eurorings Backbone AS	London UK	195.66.224.54
AS786	The JANET IP Service	London UK	195.66.226.15
AS1299	TeliaNet Global Network	London UK	195.66.224.48
AS1299	TeliaNet Global Network	London UK	195.66.226.48
AS2686	AT&T Global Network Services – EMEA	London UK	195.66.224.27
AS2818	BBC Internet Services UK	London UK	195.66.226.103
AS2856	BTnet UK Regional network	London UK	195.66.226.10
AS2856	BTnet UK Regional network	London UK	195.66.226.11
AS2914	No description available	London UK	195.66.224.138
AS2914	No description available	London UK	195.66.226.138
AS2914	No description available	London UK	2001:7f8:4::b62:1
AS3257	Tiscali Intl Network	London UK	195.66.224.32
AS3257	Tiscali Intl Network	London UK	2001:7f8:4::cb9:1
AS3291	PSINet Europe	London UK	195.66.224.14
AS3291	PSINet Europe	London UK	195.66.226.14
AS3292	TDC Data Networks	London UK	195.66.224.64
AS3292	TDC Data Networks	London UK	195.66.226.64
AS3303	Swisscom Enterprise Solutions Ltd	London UK	195.66.224.110
AS3356	Level 3 Communications	London UK	195.66.226.77
AS4589	Easynet Group Plc	London UK	195.66.224.43
AS5390	Wanadoo Nederland BV Global AS	London UK	195.66.224.31
AS5400	BT European Backbone	London UK	195.66.224.108
AS5427	Primus Telecommunications GmbH Germany	London UK	195.66.224.106
AS5427	Primus Telecommunications GmbH Germany	London UK	195.66.226.106
AS5430	freenet City LINE GmbH	London UK	195.66.224.102
AS5511	France Telecom	London UK	195.66.224.83
AS5571	Netcom Internet Ltd	London UK	195.66.224.33
AS5571	Netcom Internet Ltd	London UK	195.66.226.33
AS5604	Freedom To Surf Plc	London UK	195.66.224.41
AS5669	VIA NET.WORKS Inc	London UK	195.66.226.28
AS6656	Star Internet Ltd	London UK	195.66.224.127
AS6730	sunrise (TDC Switzerland AG)	London UK	195.66.224.85
AS6779	ICLnet	London UK	195.66.226.80
AS6805	Telefonica Deutschland Autonomous System	London UK	195.66.226.57
AS6830	UPC Distribution Services	London UK	195.66.224.89
AS8272	Netscalibur UK	London UK	195.66.226.47
AS8406	PIPEX Communications	London UK	195.66.224.71
AS8406	PIPEX Communications	London UK	195.66.226.71
AS8422	NETCOLOGNE AS	London UK	195.66.224.172
AS8426	ClaraNET	London UK	195.66.226.66
AS8586	REDNET Ltd	London UK	195.66.226.73
AS9019	DATAGRAMA AS	London UK	195.66.226.149

AS12390	Kingston Communications plc AS	London UK	195.66.226.119
AS12513	Eclipse Internet	London UK	195.66.224.117
AS12621	1A Networks Limited	London UK	195.66.224.113
AS126211	A Networks Limited	London UK	195.66.226.113
AS12932	Teletext Ltd.	London UK	195.66.226.123
AS12956	Telefonica Backbone Autonomous System	London UK	195.66.224.134
AS13127	AS for the Trans-European Versatel IP Transport backbone	London UK	195.66.226.142
AS13129	Global Access Telecommunications Inc.	London UK	195.66.224.132
AS13184	HanseNet Telekommunikation GmbH	London UK	195.66.224.104
AS13237	LambdaNet AS for European Operations	London UK	195.66.224.99
AS13285	Opal Telecom	London UK	195.66.224.136
AS13646	Priority Telecom Global Autonomous System	London UK	195.66.224.118
AS13646	Priority Telecom Global Autonomous System	London UK	195.66.226.118
AS15444	Netservices Plc	London UK	195.66.224.109
AS15444	Netservices Plc	London UK	195.66.226.109
AS15861	Legend Internet	London UK	195.66.224.40
AS15861	Legend Internet	London UK	195.66.226.40
AS20500	Griffin Internet European Network	London UK	195.66.224.38
AS20679	hSo: broadband internet data and	London UK	195.66.224.160
AS20718	arsys.es	London UK	195.66.224.165

Additional lists of other RRCs are available in [29].

Appendix B List of programs and script used to preprocess data

Appendix B-1 Data filtering procedure.

1. uncompress data (if data is NOT downloaded by MSIE) by gunzip or else;
2. “./route_btoa -m -i filename > xxx” to convert data from binary format to ASCII format and then save into a file;
3. “java extract xxx xxx.log” to extract AS_PATH from the data;
4. “sed -e ‘s/[]//g’ xxx > xxxx” to delete “[” and “]” from data (unique in Ripe);
5. “java AdjMatrix “ to get the adjacent pairs (specify filenames in Java source code);
6. calculate degrees with Deg.java;
7. use “gen_L.m” in Matlab to create the adjacent matrix.

Appendix B-2 Program (Exact.java) used to extract AS_Path from data.

```
public class Extract {

    public Extract(String FileName, String Output_File) {

        try{
            BufferedReader BR_InputFile = new BufferedReader (new FileReader(FileName));
            FileWriter FW_OutputFile = new FileWriter (Output_File);
            String EachLine = null;

            while ( ( EachLine = BR_InputFile.readLine()) != null ) /*&& Index_File < 150*/ {
                StringTokenizer ST_EachLine = new StringTokenizer(EachLine, "|");
                int count = 0;
                while (ST_EachLine.hasMoreTokens() ){
                    String Eachword = (String)(ST_EachLine.nextToken());
                    if (count==6){
                        FW_OutputFile.write(Eachword);
                        FW_OutputFile.write('\n');
                        break;
                    }
                    count+=1;
                }
            }

            BR_InputFile.close();
            FW_OutputFile.close();
        }
    }
}
```

```

    }
    catch (IOException IoEx){
        IoEx.printStackTrace();
    }
}
public static void main(String[] args) {
    Extract ext = new Extract(args[0], args[1]);
}
}

```

Appendix B-3 Program (AdjMatrix.java) used to build adjacent matrix

```

public class AdjMatrix {
    final static String Null_Node = "000000";
    public AdjMatrix() {
    }
    void generateMatrix(String FileName, String Output_File){
        Vector Matrix = new Vector();
        try{
            BufferedReader BR_InputFile = new BufferedReader (new FileReader(FileName));
            FileWriter FW_OutputFile = new FileWriter (Output_File);
            String EachLine = null;
            boolean Before_Is_Null_Node = false;
            while ( ( EachLine = BR_InputFile.readLine()) != null){
                String Pre_Node = null, Curr_Node = null;
                StringTokenizer ST_EachLine = new StringTokenizer(EachLine, " ");
                int Index_EachLine = 0;
                while (ST_EachLine.hasMoreTokens()){
                    String EachWord = ST_EachLine.nextToken();

                    // the AS before exists, need to calculate connectivity.
                    if ( Pre_Node != null ){
                        int Curr_AS_No = 0;
                        int Curr_Matrix_Size = 0;
                        int Before_AS_No = 0;
                        // current AS isnt null
                        if ( ! EachWord.trim().equals(Null_Node)){
                            Curr_AS_No = Integer.parseInt(EachWord.trim());
                            Before_AS_No = Integer.parseInt(Pre_Node.trim());
                            Curr_Matrix_Size = Matrix.size();
                            Pre_Node = EachWord.trim();
                            // the current AS # exceeds the matrix
                            if ( Curr_Matrix_Size < Curr_AS_No ) {
                                int Gap = Curr_AS_No - Curr_Matrix_Size;

                                for (int i = 0; i < Curr_Matrix_Size; i++){ // expand current vectors
                                    Vector Curr_V = (Vector)Matrix.elementAt(i);
                                    for (int j = 0; j < Gap; j++)
                                        Curr_V.add("0");
                                }
                            }
                        }
                    }
                }
            }
        }
    }
}

```

```

    }

    for (int i = Curr_Matrix_Size; i < Curr_AS_No; i++){ //add new vectors
        Vector New_V = new Vector();
        for (int j = 0; j < Curr_AS_No; j++)
            New_V.add("0");
        Matrix.add(New_V);
    }

    ((Vector)Matrix.elementAt(Curr_AS_No - 1)).set(Before_AS_No - 1, "1");
    ((Vector)Matrix.elementAt(Before_AS_No - 1)).set(Curr_AS_No - 1, "1");
}
// the current AS# is included in the matrix
else{
    ((Vector)Matrix.elementAt(Curr_AS_No - 1)).set(Before_AS_No - 1, "1");
    ((Vector)Matrix.elementAt(Before_AS_No - 1)).set(Curr_AS_No - 1, "1");
}
}
// current node is null
else
    Pre_Node = null;

}
// we dont need to calculate connectivity,
// but store AS#.
else if ( ! EachWord.trim().equals(Null_Node))
    Pre_Node = EachWord.trim();

} //while
} //while
for (int i = 0; i < Matrix.size(); i++){
    Vector Row_V = (Vector)Matrix.elementAt(i);
    for (int j = 0; j < Matrix.size(); j++){
        FW_OutputFile.write( ((String)Row_V.elementAt(j)).trim()+" ");
    }
    FW_OutputFile.write("\n");
}

FW_OutputFile.close();
BR_InputFile.close();
}
catch (IOException IoEx){
    IoEx.printStackTrace();
}
}
/**Main method*/
public static void main(String[] args) {
    AdjMatrix AdjM = new AdjMatrix();
    AdjM.generateMatrix("ASoutput.log", "RouteViews_AdjMatrix.log");
}
}

```

Appendix B-4 Program (Deg.java) used to calculate degrees of ASs.

```
public class Deg {

    public Deg(String FileName, String Output_File) {

        int degree = 10000;
        try{
            BufferedReader BR_InputFile = new BufferedReader (new FileReader(FileName));
            FileWriter FW_OutputFile = new FileWriter (Output_File);
            String EachLine = null;
            int Degrees[] = new int [degree];
            for (int i = 0; i<degree; i++)
                Degrees[i]=0;
            while ( ( (EachLine = BR_InputFile.readLine()) != null) /*&& Index_File < 150*/ ){
                StringTokenizer ST_EachLine = new StringTokenizer(EachLine);
                while (ST_EachLine.hasMoreTokens() ){
                    String EachWord = ST_EachLine.nextToken();
                    if (Integer.parseInt(EachWord) <= 10000)
                        Degrees[Integer.parseInt(EachWord)-1] = Degrees[Integer.parseInt(EachWord)-1]
+ 1;
                }//while
            }
            for (int i = 0; i<degree; i++){
                FW_OutputFile.write(String.valueOf(Degrees[i]));
                FW_OutputFile.write('\n');
            }
            FW_OutputFile.close();
            BR_InputFile.close();
        }
        catch (IOException IoEx){
            IoEx.printStackTrace();
        }
    }

    public static void main(String[] args) {
        Deg deg = new Deg(args[0], args[1]);
    }
}
```

Appendix B-5 Program (gen_l.m) used to generate the Normalized Laplacian matrix.

```
%generate Normalized Laplacian
caida_L3=sparse(30000,30000);
for m=1:30000
    if (all_degrees(m,1)~=0)
        caida_L3(m, m)=1;
```

```

    end
end
fprintf('diagonal done!\n')
for m=1:18837
    if ((all_adj(m,1)<=30000) && (all_adj(m,2)<=30000))
        %if ((all_degrees(all_adj(m,1),1)~=0) && (all_degrees(m,2)~=0))
        temp = -(1/(sqrt(all_degrees(all_adj(m,1),1)*all_degrees(all_adj(m,2),1)) ));
        caida_L3(all_adj(m,1), all_adj(m,2))= temp;
        caida_L3(all_adj(m,2), all_adj(m,1))= temp;
    end
end;
fprintf('all done!')

```

References

- [1] AS paths to individual networks: <http://moat.nlanr.net/ASx/> (15, Aug. 2004).
- [2] Autonomous System Numbers: <http://www.iana.org/assignments/as-numbers> (15, Aug. 2004).
- [3] A. Barabási and R. Albert, "Emergence of scaling in random networks," *Science*, vol. 286, pp. 509-512, Oct. 1999.
- [4] BGP-system usage of 32 bit Internet address space: <http://moat.nlanr.net/IPaddrocc/> (15, Aug. 2004).
- [5] Y. Rekhter and T. Li, "A Border Gateway Protocol 4 (BGP-4)," RFC 1771, March 1995.
- [6] BGP Reports: <http://www.potaroo.net/bgp/> (15, Aug. 2004).
- [7] K. Calvert, M. Doar, and E. W. Zegura, "Modeling Internet topology," *IEEE Communications Magazine*, vol. 35, no. 6, June 1997, pp. 160–163.
- [8] J. M. Carlson and J. Doyle, "Highly optimized tolerance: a mechanism for power laws in designed systems," *Physical Review Letter*, Apr. 1999, vol. 60, no. 2, pp. 1412–1427.
- [9] H. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger, "Towards capturing representative AS-level Internet topologies," *Computer Networks Journal*, vol. 44, no. 6, pp. 735-755, Apr. 2004.
- [10] H. Chang, S. Jamin, and W. Willinger, "What causal forces shape Internet connectivity at the AS-level?" *Technical report, EECS Dept., Univ. of Michigan*, 2003.
- [11] J. Chen and Lj. Trajković, "Analysis of Internet topology data," *IEEE Int. Symp. Circuits and Systems*, Vancouver, British Columbia, May 2004, vol. IV, pp. 629–632.
- [12] F. R. K. Chung, *Spectral Graph Theory*. Providence, Rhode Island: Conference Board of the Mathematical Sciences, 1997, pp. 2–6.
- [13] M. Doar, "A better model for generating test networks," *Proc. of GLOBECOM'96*, London, UK, Nov. 1996, pp. 86–93.
- [14] J. Doyle and J. M. Carlson, "Power laws, highly optimized tolerance and generalized source coding," *Physical Review Letter*, vol. 84, no. 24, pp. 5656–5659, Mar. 2000.
- [15] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On power-law relationships of the Internet topology," *Proc. of ACM SIGCOMM '99*, Cambridge, MA, Aug. 1999, pp. 251–262.
- [16] M. Fiedler, "A property of eigenvectors of nonnegative symmetric matrices and its application to graph theory," *Czech Math Journal*, vol. 25, no. 1, pp. 619–633, 1975.

- [17] M. Fiedler, "Algebraic connectivity of graphs," *Czech Math Journal*, vol. 23, no. 2, pp. 298–305, 1973.
- [18] L. Gao, "On inferring autonomous system relationships in the Internet," *IEEE/ACM Transactions on Networking*, Dec. 2001, vol. 9, no. 6, pp. 733–745.
- [19] L. Gao, T. G. Griffin, and J. Rexford, "Inherently safe backup routing with BGP," *Proc. of Infocom 2001*, Anchorage, Alaska, April, 2001, vol. 1, pp. 547–556.
- [20] L. Gao and J. Rexford, "Stable Internet routing without global coordination," *Proc of ACM SIGMETRICS*, Santa Clara, CA, June 2000, pp. 307–317.
- [21] G. Huston, "Interconnection, peering and settlements-Part I," *Internet Protocol Journal*, Mar. 1999: http://www.cisco.com/warp/public/759/ipj_2-1/ipj_2-1_ps1.html (15, Aug. 2004).
- [22] G. Huston, "Interconnection, peering and settlements-Part II," *Internet Protocol Journal*, June 1999: http://www.cisco.com/warp/public/759/ipj_2-2/ipj_2-2_ps1.html (15, Aug. 2004).
- [23] A. Lakhina, J. W. Byers, M. Crovella, and I. Matta, "On the geographic location of Internet resources," *Proc. of 2nd ACM SIGCOMM Internet Measurement Workshop*, Marseille, France, Nov. 2002, pp. 249–250.
- [24] Looking Glass sites: <http://www.traceroute.org> (15, Aug. 2004).
- [25] A. Medina, I. Matta, and J. W. Byers, "On the origin of power laws in Internet topologies," *ACM SIGCOMM Computer Communication Review*, vol. 30, no. 2, pp. 18–28, Apr. 2000.
- [26] M. Mihail, C. Gkantsidis, and E. Zegura, "Spectral analysis of Internet topologies," *Proc. of Infocom 2003*, San Francisco, CA, Mar. 2003, vol. 1, pp. 364–374.
- [27] National Laboratory for Applied Network Research: <http://nlanr.net/> (15, Aug. 2004).
- [28] Netgeo: <http://www.caida.org/tools/utilities/netgeo/> (15, Aug. 2004).
- [29] Réseaux IP Européens: <http://www.ripe.net/ris> (15, Aug. 2004).
- [30] E. Chen and J. Stewart, "A Framework for Inter-Domain Route Aggregation," RFC 2519, February 1999.
- [31] Route Views project: <http://www.routeviews.org> (15, Aug. 2004).
- [32] Skitter: <http://www.caida.org/tools/measurement/skitter/> (15, Aug. 2004).
- [33] L. Subramanian, V. N. Padmanabhan, and R. H. Katz, "Geographic properties of Internet routing," *Proc. of USENIX Annual Technical Conference*, Monterey, CA, Jun. 2002, pp. 243–259.
- [34] H. Tangmunarunkit, J. Doyle, R. Govindan, S. Jamin, and S. Shenker, "Does AS size determine degree in AS topology?," *ACM SIGCOMM Computer Communication Review*, Oct. 2001, vol. 31, no. 5, pp. 7–10.
- [35] The Cooperative Association for Internet Data Analysis: <http://www.caida.org> (15, Aug. 2004).
- [36] The SCAN project: <http://www.isi.edu/scan/> (15, Aug. 2004).

- [37] D. Vukadinovic, P. Huang, and T. Erlebach, "On the spectrum and structure of Internet topology graphs," in H. Unger et al., Editors, *Innovative Internet Computing Systems*, LNCS2346, Springer, Berlin, Germany, 2002, pp. 83–96.
- [38] F. Wang and L. Gao, "On inferring and characterizing Internet routing policies," *Proc. of ACM SIGCOMM Internet Measurement Conference*, Miami Beach, FL, Oct. 2003, pp. 15–26.
- [39] B. M. Waxman, "Routing of multipoint connections," *IEEE Journal on Selected Areas in Communications*, vol. 6, no. 9, pp. 1617–1622, Dec. 1988.
- [40] X. Zhao, D. Pei, L. Wang, D. Massey, A. Mankin, S. Wu, and L. Zhang, "An analysis of BGP multiple origin AS (MOAS) conflicts," *Proc. ACM SIGCOMM Internet Measurement Workshop*, San Francisco, CA, Nov. 2001, pp. 31–35.