DETERMINING THE LOCATION OF AUTONOMOUS SYSTEM RELATIONS AMONG TOP TIER INTERNET SERVICE PROVIDERS IN THE UNITED STATES

by

PHILLIP KRIEGEL

A THESIS

Presented to the Department of Computer and Information Science
and the Robert D. Clark Honors College
in partial fulfillment of the requirements for the degree of
Bachelor of Science

June 2016

An Abstract of the Thesis of

Phillip Kriegel for the degree of Bachelor of Science
in the Department of Computer and Information Science to be taken June 2016

Title: Determining the Location of Autonomous System Relations Among Top Tier Internet Service Providers in the United States

Approved: _______________

Professor Reza Rejaie

The Internet is a network of networks. Understanding where and how these networks interconnect is important for the purpose of meaningfully investigating a wide range of critical Internet-related problems, such as the vulnerability of the Internet to physical damage. While there is published work on Internet topology and structure, those studies focus primarily on finding the existence of network interconnections and characterizing the structure of the Internet by those relationships. What this thesis seeks to investigate is a methodology that helps to determine where these interconnections, or 'cross connects', are happening at the city-level. We evaluate a method for collecting cross connect data that uses geo-located vantage points and targets. We then investigate the feasibility of using distributions of round trip times to estimate a threshold for separating cross connects that occurred within a city from those that occurred outside of a city. We utilized this method to investigate the cross connects in 17 cities across the United States. Preliminary findings indicate that this method is viable in certain cities. The study also identified several trends that warrant further investigation.

Acknowledgements

This thesis in many ways represents the capstone to my undergraduate studies. As such, I would like to thank everyone who helped me not only work on this thesis, but those who have supported me throughout my time at the University of Oregon.

To Professor Rejaie, Reza Motamedi, and the rest of ONRG: thank you for your support, for accepting me into your group, and for your unbelievable patience. I really didn’t deserve it. It has been an absolute pleasure working with you.

To my friends: I chose to surround myself with you because I knew you would bring out the best in me. I am privileged to have had you as companions for this stage of my life, and look forward to having you as friends for the rest of my life. I know that I would not be where I am now without you.

To my parents: excellence was never an option, but an expectation. You never asked me to do my best; you demanded it. My will, my drive, and my motivation come directly from you. You have been the root of all of my success and I can only hope that I have made you proud.

Table of Contents

Chapter 1: Introduction 1

1.1 - Problem Statement and Preview of Approach 2

1.2 - Scope and Limitations 3

1.3 - Key Findings and Contributions 4

1.4 - Overview of Remaining Chapters 4

Chapter 2: Background 6

2.1 - Internet Connectivity and Routing 6

2.2 - Routing 8

2.3 - Tools Used for Measurement 9

2.4 - Online Information 11

2.5 - Measurement Platform 13

2.6 - Glossary of Terms 13

Chapter 3: Methodology 16

3.1 - Summary of Methodology 16

3.2 - Identifying the Vantage Points and Targets for Traceroute Measurements 16

3.3 - Conducting the Measurements and Collecting the Data 18

3.4 - Inferring Cross Connect Location Using Traceroute Measurement Data 20

3.5 - Validating Inferences with an External Source 20

3.6 - Data Analysis 21

Chapter 4: Data Collection and Data Set 23

4.1 - Experimental Setting 23

4.2 - Traceroute Dataset 26

4.3 - Ping Dataset 30

4.4 - PeeringDB Data Set 31

4.5 - RTT Threshold Data Set 31

Chapter 5: City-Level Analysis 37

5.1 - Number of Cross Connects Found & Evidence of AS Presence 37

5.2 - How Many AS-Pairs Present in a City According to Peering DB Have Established a Cross Connect? 44
5.3 - Which ASes Have the Largest Presence in a City? 47
Chapter 6: AS-Level Analysis 57
6.1 - In Which Cities Does an AS Have the Largest/Smallest Presence? 57
6.2 - Which ASes Are Present in the Most Cities? 61
6.3 - Which ASes Predominantly Serve as Facilitators, and Which Are Recipients? 63
Chapter 7: AS-Pair-Level Analysis 66
7.1 - Do ASes Favor Some Cities More in Terms of the Number of Cross Connects That They Have Present in a City? 66
7.2 - Is There a Pattern Regarding Where Cross Connects Between AS-Pairs Are Located? 68
7.3 - Of the ASes Present in a City, What Percentage of Possible AS-Pairs Have a Cross Connect Between Them? 69
Chapter 8: Related Work 72
Chapter 9: Conclusions and Future Work 74
Chapter 10: References 77

List of Figures

Figure 2.1: A sample diagram depicting the geo-footprint of AS3356, AS701, and AS209. The vertical lines extend to major cities within the footprint [6]........................ 7

Figure 2.2: A router path traversed by a traceroute, which represents it as a path of IP interfaces [4].. 9

Figure 2.3: An examp le of an IP path mapped to an AS path. The connections between ASes are cross connects. They can either occur within a city or outside of a city.......... 10

Figure 4.1: Probe coverage images (courtesy of RIPE Atlas). Each dot represents a probe. Green dots represent connected probes, yellow dots represent disconnected probes, and red dots indicate abandoned probes. Starting clockwise from the top left: Dallas, Chicago, New York, and Los Angeles... 25

Figure 4.2: Tables of traceroute validation statistics for Chicago, Dallas, Los Angeles, and New York.. 28

Figure 4.3: RTT CDF of both traceroute and ping measurements in New York City ... 32

Figure 4.4: RTT CDF of both traceroute and ping measurements in Los Angeles........ 33

Figure 4.5: RTT CDF of both traceroute and ping measurements in Dallas.................. 34

Figure 4.6: RTT CDF of both traceroute and ping measurements in Dallas.................. 35

Figure 5.1: Matrix of detection and corroboration of cross connects in Los Angeles ... 39

Figure 5.2: Matrix of detection and corroboration of cross connects in Dallas 40

Figure 5.3: Matrix of detection and corroboration of cross connects in Chicago 41

Figure 5.4: Matrix of detection and corroboration of cross connects in New York City43

Figure 5.5: Bar chart of AS involvement in cross connects in Los Angeles.................. 48

Figure 5.6: CDF of cross connects found for each AS-pair in Los Angeles 49

Figure 5.7: Bar chart of AS involvement in cross connects in Dallas............................ 50

Figure 5.8: CDF of cross connects found for each AS-pair in Dallas 51

Figure 5.9: Bar chart of AS involvement in cross connects in Chicago......................... 52

Figure 5.10: CDF of cross connects found for each AS-pair in Chicago....................... 53

Figure 5.11: Bar chart of AS involvement in cross connects in New York City 54

Figure 5.12: CDF of cross connects found for each AS-pair in New York City 55

Figure 6.1: Bar chart of cross connects found in each city for AS3356......................... 58

Figure 6.2: Bar chart of the number of ASes found present in each city 59

Figure 6.3: Bar chart of the number of cities each AS was present in 62

Figure 6.4: Bar chart of the number of cross connects AS3356 has with top 20 ASes.. 64

Figure 6.5: Bar chart of the number of cross connects AS6453 has with top 20 ASes.. 64

Figure 7.1: Bar chart showing number of cities each AS was present in on top of a scatter plot of cross connects per city per AS 67
Figure 7.2: CDF showing the number of cross connects found for each AS-pair 70
Figure 7.3: CDF showing the number of cities each AS-pair was present in 70

List of Tables

Table 4.1.: Targeted cities and their RIPE Atlas probe count 23

Table 4.2: Selected ASes 25

Table 4.3: Number of traceroutes launched for each city 29

Table 4.4: Number of pings launched for each city 30

Table 4.5: Estimated RTT threshold for each city 36

Table 5.1: Cross connect evidence summary for Los Angeles 39

Table 5.2: Cross connect evidence summary for Dallas 40

Table 5.3: Cross connect evidence summary for Chicago 41

Table 5.4: Cross connect evidence summary for New York City 43

Table 5.5: PeeringDB evidence summary for Los Angeles, Dallas, Chicago, and New York 45

Table 7.1: City groups with the most common cross connect counts for groups of cities size 2 to 9 68

Chapter 1: Introduction

The Internet is a network of networks. There are more than 40,000 networks that combine to form today's Internet. A network will transmit packets of data between routers contained within itself, as well as to routers at the edge of other networks, in order to move information from one machine to another. Accessing a website, watching an Internet video, or posting a status to social media begins a process that transfers data across dozens of routers over thousands of miles in mere milliseconds. Most networks are managed by a company or an organization, such as AT&T, Comcast, or a university. This sort of a network is known as an autonomous system, or AS. These networks work together to allow access and communication between users around the world.

Understanding the topology of the Internet is of great interest, as it reveals vulnerabilities of the Internet to various attacks. Prior studies have either focused on logical relationships between different networks. Others have tried to capture the topology at the router-level. This latter form of topological study is extremely difficult due to the number of routers present in the Internet, as well as its constantly evolving nature.

Therefore, this thesis chooses to focus on determining the location of an aspect of the Internet that does not change nearly as frequently: AS relationships. As mentioned before, the logical relationships are understood. It is possible to know that AS X connects with AS Y. But where does it connect? In how many places does it connect? What is the set of ASes that AS X connects with in a particular city?

Answering these questions permits a host of other questions to be answered regarding the vulnerability of the Internet to physical damage. If a city is affected by a man-made or a natural disaster, how does this disrupt or affect the Internet and the ability to access other ASes?

However, finding the cross connect between two ASes in a particular city is challenging, as it requires careful measurements that traverse these cross connects to reveal their presence. By the very nature of the Internet and network routing, it is impossible to ensure that a measurement passes through a certain geographical area, so there must be planning to increase the likelihood that a measurement passes through a city. Even then, there is an element of chance as to whether or not the measurement will reveal any relevant cross connects.

1.1 - Problem Statement and Preview of Approach

This thesis ultimately tackles two basic questions:

1. In how many cities does an AS have a cross connect with other ASes?
2. What is the set of ASes that an AS has a cross connect with within a city?

In order to answer these questions, we used a recently developed measurement technique to reveal AS cross connects in a given city [1]. The basic idea was to use traceroute measurements sent between sources and targets that are close to the center of a city. This significantly increases the chance that any physical link that is present between ASes in a city is found.

We conducted a large-scale measurement campaign based on this methodology that focused on 20 ASes and 17 cities. We utilized a largely developed platform known as RIPE Atlas to conduct our campaign, and used the results to infer cross connects between the selected ASes in each city.

We then conducted a series of analyses on the data that allowed us to understand the data on a city, AS, and AS-pair-level. By viewing the data in these different contexts, it was possibly to more fully understand the data set and what it revealed.

1.2 - Scope and Limitations

As mentioned earlier, there are over 40,000 interconnected ASes that comprise the Internet. To examine all of their interconnections would be an enormous undertaking. For this thesis, the focus was on cross connects between the top 20 largest ASes in the world as ranked by the Center for Applied Internet Data Analysis (CAIDA) [2]. These ASes are ranked by the number of ASes that they service. By using the top 20 ASes, the methodology for data collection and analysis could be developed and evaluated that could later be applied to any set of ASes, while still generating a useful data set that provided a geographical understanding of the top level AS cross connects in the United States.

Furthermore, we were constrained by the limitations of geography and RIPE Atlas probe coverage. RIPE Atlas depends on users to request probes and install them in their homes or institutions. Therefore, some cities that are excellent candidates for study based on their geography could not be investigated since they simply lack sufficient support for the RIPE Atlas platform. Another constraint of the RIPE Atlas platform is the need to gather sufficient credits to conduct measurements. Therefore, sufficient time was needed to accumulate the credits necessary to launch the campaigns.

1.3 - Key Findings and Contributions

Our key findings can be summarized as follows. For certain cities, our measurement technique returned useful results that were corroborated by external sources. The conditions and factors that caused these particular cities to be so well suited to our technique are currently unknown. Furthermore, it was found that top 20 ASes appear to underutilize the possible cross connects available to them within a city through colocation facilities and IXPs, which are buildings where many ASes converge and exchange data. Finally, it was found that the data collected could be used to infer the nature of AS relationships within a city, such as which ASes predominantly sent data to other ASes, and which ASes generally received data.

1.4 - Overview of Remaining Chapters

Chapter 2 will provide additional background and context to this thesis, while defining key terms and concepts.

Chapter 3 will cover the methodology used for this thesis. It will provide detailed descriptions of the data collection platform used for this thesis, as well the criteria for selecting ASes, targets, vantage points, and cities.

Chapter 4 will cover the data set collected for this thesis. It will provide information on the cities used, the ASes selected, and the measurement platform. It will also present characterizations and validations of the data set.

Chapter 5 will provide analysis of the data set on a city-level. It will include observations and relevant statistical analyses that were drawn from the data set for four different cities selected based upon their geographical diversity. The cities used are Los

Angeles, New York City, Dallas, and Chicago. Each form of analysis will compare and contrast the results for these four cities.

Chapter 6 will provide an analysis at the AS-level based on relevant observations found within the data. Rather than looking at each city, this analysis will encompass all of the cities that were investigated, and look at where ASes are present and to with which other ASes they form cross connects.

Chapter 7 will provide an analysis at the AS-pair-level by examining the presence of AS-pairs in cities, the number of AS-pairs that exhibited a cross connect, and sets of cities that contain cross connects between a common set of AS-pairs.

Chapter 8 will discuss related works upon which this thesis builds, and will provide commentary as to where this thesis fits within the overall body of work on this subject.

Chapter 9 will present our conclusions based on the analysis conducted, as well as provide insight into the applications of these conclusions and potential future work that could be conducted to continue the exploration of this subject.

Chapter 10 will contain references to resources used throughout this work.

Chapter 2: Background

The following chapter contains relevant background information necessary to understand the content of this thesis. Terms that are defined in this background will be aggregated at the end of this chapter in a glossary for reference.

2.1 - Internet Connectivity and Routing

The Internet is a network of interconnected networks. These networks are organized into *autonomous systems*, known as AS. An AS is a collection of connected *Internet Protocol (IP) addresses* that are under the control of an entity, such as an institution or a company [3]. An IP address is much like a street address: it allows for us to send information to a specific location. IP addresses are assigned to routers, which are also owned by the AS and facilitate the transfer of data. In the most basic sense, routers are hardware devices that send an information packet where it needs to go in order to reach its final destination. These networks of routers transfer data across the *network core* between end points on the *network edge*, such as a personal computer or a mobile phone.

ASes are tiered, with the tiers separating ASes based upon the nature of the relationships they have with each other. If an AS wants to send its network traffic through an AS on a higher tier, there is a customer-provider relationship, and the lower tier AS must pay the higher tier AS. Two ASes on the same tier are considered peers and have a peering relationship. Traffic is exchanged without any settlement [4]. Top tier ASes are ASes that participate within the Internet using only peering relationships. These ASes are so large that they can reach any other AS on the Internet without having to pay any settlements [4]. These tier 1 ASes include Internet Service Providers (ISPs) such AT&T, Sprint, Qwest, and Level3. These tiers are not strictly defined. Rather, they arose organically [5].

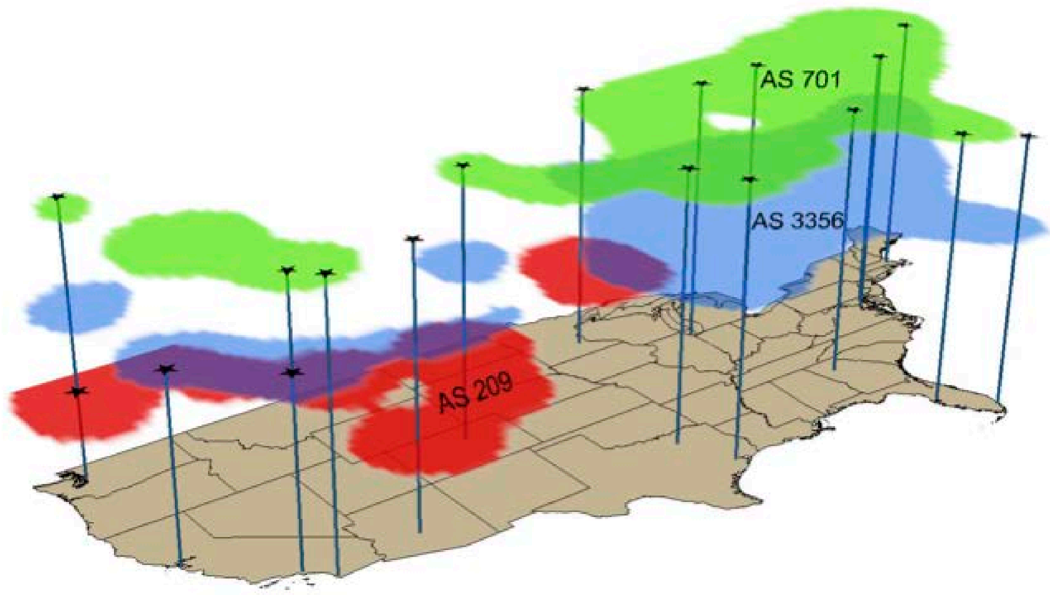

Figure 2.1: A sample diagram depicting the geo-footprint of AS3356, AS701, and AS209. The vertical lines extend to major cities within the footprint [6]

Another important aspect of internet connectivity that must be discussed are *Internet eXchange Points* (IXPs) and *colocation facilities* (also known as colos). These are physical infrastructures created and owned by an organization in order to facilitate the direct exchange of data between ASes. While ASes can create direct connections between each other, this requires the creation of infrastructure to support it. Colos and IXPs provide a central hub and infrastructure for many networks to meet and connect. This is beneficial because it reduces latency and is more cost effective [1]. The main difference between IXPs and colos is the type of connection that they facilitate. IXPs facilitate public peering, where ASes connect to each other via a shared medium such as an Ethernet switch. Colos actually create direct connections between ASes just as they could outside of a colocation facility, but the infrastructure is built and maintained by the owner of the facility.

2.2 - Routing

An extremely important concept to understand regarding the behavior of the Internet is *routing*. Every router in a network performs routing algorithms periodically to decide the best route for data. These algorithms create routing tables that let the router know where to direct data based solely upon its assigned IP address. Routers send announcements to each other that communicate information about the optimal routes to use.

Within an AS, routers utilize intra-AS routing. To send information between ASes, inter-AS routing must be used, where competing ASes send announcements across the Internet to let other ASes know about what can be reached through them. These announcements are sent out constantly, and the routing algorithms must continually be updating to account for new information. Because of this, two packets of data sent from the same source to the same destination might follow a different path, since new information might be received that indicates the superiority of another route.

Furthermore, it is important to understand that routing does not in any way reflect the physical geography of the network. Routers view the network as IP addresses that are a certain number of milliseconds of delay away. This is almost entirely independent of geography due to the speed of electronic data transfer. While distance between routers can certainly affect the delay between them, factors that are more likely to influence delay on the network include congestion, queuing delays, and load balancing [1]. Because of this, two users in the same city communicating with each other via the Internet might have their messages pass through several other cities before arriving to their intended destination.

2.3 - Tools Used for Measurement

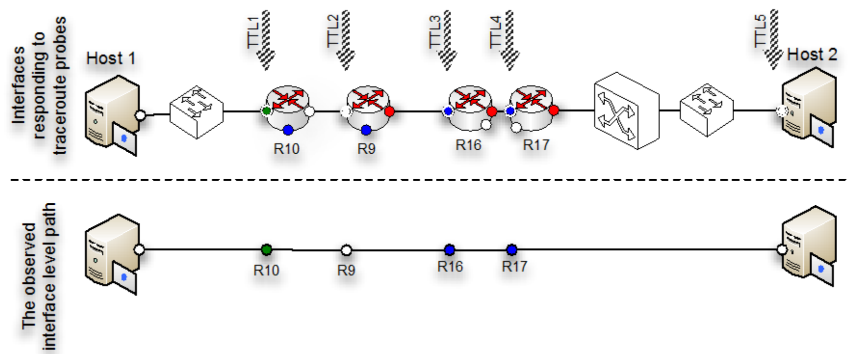

Figure 2.2: A router path traversed by a traceroute, which represents it as a path of IP interfaces [4]

Traceroute and Ping

An important tool used for this research is the *traceroute* program. A traceroute measurement is launched from a source host, or *vantage point*, to a remote destination, or *target*. Along the way, the measurement reports back the IP addresses of the router interfaces that it passed through. Using this tool, it is possible to have an IP address-level view of the path that a packet takes from a vantage point to a target, as seen in figure 2.1 [4]. A traceroute also returns the *round trip time (RTT)* for each router along the path. The RTT is the amount of time it took for the measurement to reach the router and return to the source.

It is important to note that the results from running a traceroute with the same vantage point and target multiple times can differ. As discussed earlier, routers might direct different packets of data along different paths based on a variety of factors, such as load balancing on the network to prevent overwhelming any one router with traffic. Even after sending multiple traceroutes from a vantage point to a destination, it is not possible to know if every potential path or router between the endpoints has been traversed.

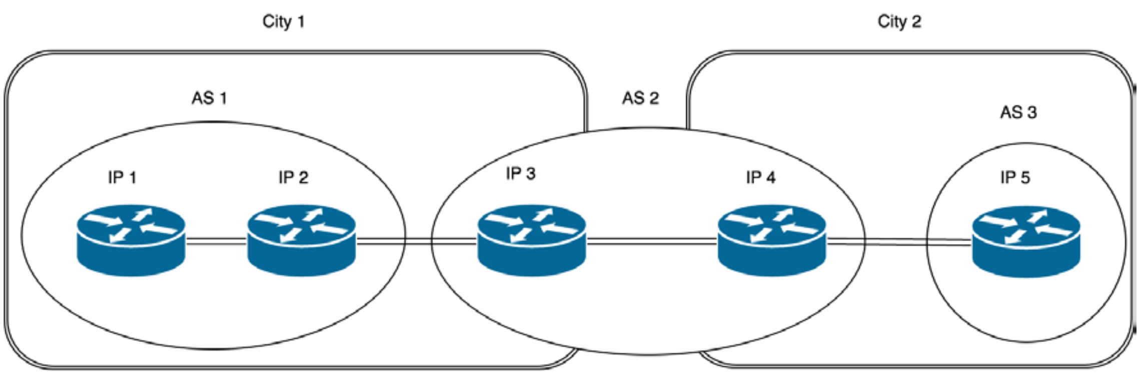

Figure 2.3: An examp le of an IP path mapped to an AS path. The connections between ASes are cross connects. They can either occur within a city or outside of a city.

As discussed earlier, an IP address belongs to an AS. Using IP-to-AS resolution tools such as Whois, which is developed by Team Cymru, it is possible to convert an IP address level view of a path to an AS-level view [7]. Once this is accomplished, one must simply examine the path to find where one AS changes to another. This is a cross connect between two ASes, and indicates that the two ASes connect either through a direct link managed by the ASes themselves, or through a colo or IXP.

While IP addresses in a traceroute can be mapped to an AS, it is not so simple to map an IP address to a geographical location [1, 8, 6]. IP addresses do not implicitly map to a geographic location, and most IP geolocation services, such as IP2Location and MaxMind [9, 10] can only provide a rough estimate. This is made even more difficult by the fact that geolocation of routers in the core is much less accurate that the geolocation of routers on the edge [11]. Since this study is looking for cross connects, the routers that will be investigated will predominantly be in the core of the network. In order to increase the likelihood that a traceroute measurement passes through a certain area, the vantage point and target of the measurement can be selected to be very close to each other and within the proximity of a well-known hub, such as a large city. This increases the probability that the traceroute measurement will pass through the targeted city [8].

Another measurement tool that we used was the *ping* program. A ping measurement is very similar to a traceroute measurement in that it sends a packet to a target from a vantage point in order to gather data regarding the network. However, unlike a traceroute measurement, a ping measurement simply returns the RTT that it took to get to the target. This made it ideal for situations that called for quick data collection that was light on resources and extraneous information.

2.4 - Online Information

There are several resources found online that were used in this thesis. These assisted in determining targets for measurements, resolving IP addresses to ASes, selecting ASes, and getting information about IXPs and colos.

The first of these is Akamai EdgeScape [12]. This is a service that is capable of providing geolocation estimates for the targets used in this thesis. In order to make it more likely that a traceroute passed through a city, a source had to be selected near the city, and a target had to be selected near the city. Using EdgeScape, it was possible to select IP addresses on the edge of the network that could be targeted by traceroute measurements. EdgeScape provides geolocation estimates, but they are only considered accurate for devices on the edge of the network, such as personal computers and servers [11]. EdgeScape is not accurate enough to geo-locate routers in the network core.

Another resource was Team Cymru's Whois service. As mentioned earlier, this service allows us to resolve IP addresses to their AS, which enables an AS-level view of a traceroute path [7].

The next resource used was the Center for Applied Internet Data Analysis (CAIDA), and in particular, their AS rankings. This ranks the top 20 ASes in the world according to the number of customers that they serve. Using this resource, the 20 ASes that were used in this thesis were selected.

Finally, this thesis leverages the information stored in PeeringDB, a service that provides information about the ASes that are located in IXPs and colos. PeeringDB was used to find AS-pairs that were present in the same city. PeeringDB can only state whether or not an AS is present at a facility; it cannot provide any information regarding which ASes that AS has chosen to connect with. Thus, the presence of two AS in a facility only suggests that they might have a cross connect there. This was used primarily to provide additional evidence that our method was working, while also allowing us to see whether or not ASes were utilizing the opportunities they had to connect with other ASes within IXPs and colos [6, 13].

2.5 - Measurement Platform

The traceroute measurements that were launched for this thesis had to be launched from a known geographical location in order to increase the probability of the measurement passing through the area of the city. As mentioned earlier, targets were selected using EdgeScape. Sources were selected and traceroutes were launched using a platform called *RIPE Atlas.* Maintained by the Réseaux IP Européens (RIPE) Network Coordination Center, this Internet measurement network spans over 9000 probes across the world. Using these probes, a member can launch various measurements, such as ping and traceroute. Since each user registers his or her probe with a physical street address as well an IP address, these probes can be selected based on their geographical location as well as on the AS to which they pertain [14]. To become a member, one must request a probe from RIPE. Upon installation of the probe, it becomes available for use by other members around the world. Every hour of uptime is rewarded with credits, which the user can then use to launch his or her own measurements from any probe in the world.

Using this measurement platform, it was possible to select a source near the center of a city. Once EdgeScape provided a target near the city, RIPE Atlas allowed us to launch a traceroute or ping measurement towards that IP address.

2.6 - Glossary of Terms

Internet Protocol (IP) Address: A string of four numbers that is used to direct information to the correct user or machine.

Autonomous System (AS): A collection of connected IP addresses that are under the control of an entity, such as an institution or a company.

Router: Hardware that performs routing algorithms to direct information packets to the next router based on the IP address assigned to the packet.

Routing: The process of determining where a packet should be sent in such a way that the network is able to run more efficiently. This means understanding the fastest way to get to a specific destination while also performing load balancing to keep the network from becoming congested.

Round Trip Times (RTT): The amount of time necessary for a packet of data to travel from the source to the destination and back.

Network Edge: The part of the Internet that faces the user. Phones, laptops, etc.

Network Core: The part of the Internet that users do not see, which generally consists of routers.

Internet eXchange Point (IXP): A facility where ASes can become tenants and engage in public peering with other ASes.

Colocation Facility (colo): A facility where ASes can become tenants and engage in private peering with other ASes.

Traceroute: A program that finds the IP level representation of a path between two end points on the network. Also reports RTT from the source to each router along the way.

Ping: A program that simply provides the RTT between the source and the destination.

Vantage Point/Target: A vantage point is a probe or a host from which measurements can be launched towards targets. Every measurement needs a well-defined vantage point and a well-defined target.

Cross Connect: The connection between two routers that are each in different ASes.

RIPE Atlas: A distributed Internet measurement platform that allows users to launch network measurements from any of over 9000 probes scattered across the world.

Chapter 3: Methodology

The following methodology will be expressed in two separate formats. First, a summary of our methodology will be shown that provides a high level view of the steps involved. This will be followed by a series of subsections that describe each major step in detail.

3.1 - Overview of Methodology

This thesis involved the utilization of traceroute and ping measurements sent between vantage points and target IP addresses that had been specifically selected to maximize the chances of the measurement passing through the targeted city. Using RIPE Atlas, probes were selected as vantage points within the vicinity of the city. EdgeScape was used to select targets within the same area. By sending traceroute and ping measurements between these vantage points and targets, there was a chance that a cross connect between ASes of interest would be discovered. The data collection campaign is inherently opportunistic: there is a chance that a cross connect might not be found, and if it is, there might be indicators that the cross connect did not occur within the city. Two indicators qwew used to determine the degree of confidence with which it can be said that a cross connect did occur within the city limits: an RTT threshold estimated for each city, and PeeringDB knowledge of AS presence in IXPs and colos.

3.2 - Identifying the Vantage Points and Targets for Traceroute Measurements

City Selection

Since the number of traceroute measurements that could be performed was limited, it was decided to focus on a subset of American cities and perform more robust analyses of each one. By looking at major AS network maps and cross referencing them with RIPE Atlas' probe coverage maps, it was possible to see which cities served as network hubs while also having enough probes to make sufficient data collection possible. In order to qualify, a city had to have at least 15 probes within a 100-mile radius of the its geographical center.

AS Selection

For this thesis, the top 20 ASes (as ranked by CAIDA based on the number of customers that they service) were selected for investigation. The only cross connects that were sought after involved these 20 ASes. The top 20 ASes can be viewed in table 4.1.

Vantage Point Selection

Once the cities and ASes were selected, vantage points were selected in each city using RIPE Atlas. For this thesis, RIPE Atlas probes were used as vantage points. All of the probes within a 100-mile radius of the center of a given city were selected for use as a vantage point.

Destination Selection

In addition to a vantage point, a traceroute or ping measurement needs a destination. Using EdgeScape, it was possible to select destinations within a 100-mile radius of the city center. Up to three destinations were chosen per top 20 AS based on two main factors. The first was the target's proximity to the center of the city. The second was the uniqueness of the target's subnet. This was done in order to increase the likelihood of a successful traceroute. Unique prefixes were preferred because an issue encountered by a traceroute to an IP address in a particular subnet might be replicated if a backup went to a different IP address in the same subnet. Thus, three IP address targets in three different subnets within the same AS were used whenever possible.

3.3 - Conducting the Measurements and Collecting the Data

Campaign Launch

Once all of the destinations and vantage points were gathered for a city, traceroute measurements were launched from every vantage point to every destination using RIPE Atlas. Ping measurements were launched as well between all pairs of the selected vantage points. Since up to 3 targets were selected per top 20 AS, there was a maximum of 60 targets per campaign. The number of ping measurements launched was the number of possible pairs between each of the N probes, which is given by the formula N(N-1)/2.

Traceroute measurements cost on average 40 RIPE Atlas credits, and ping measurements cost 3 RIPE Atlas credits. In total, nearly half a million credits were spent for all of the campaigns. To accumulate that amount of credit required about one month of hosting our own RIPE Atlas probe, which made extensive testing prior to the launch of the entire campaign of the utmost importance. With a single press of a button, a month’s worth of resources could have been wasted.

Data Collection

A full 60 minutes was allowed to pass after the campaign launch to ensure that all traceroute and ping measurements would have either completed or failed. The data was then collected and placed into a MySQL database. Each city was assigned its own database schema with several tables. These tables stored relevant data about the traceroute and ping measurements, such as the intended target, whether or not it succeeded, and the vantage point from which it was launched. There was also a table that included every single hop taken by the traceroute measurements. The ping measurements were also collected, and their average RTT was stored in the database.

Data Processing and Preparation

The next step was to take the data as it existed and to process it to be more useful. Each traceroute hop's IP addresses were mapped to their respective ASes using Team Cymru's Whois service. After this, it was possible to see when a traceroute transitioned from one AS to another. A table was also created that contained every single cross connect between two different ASes, and specified exactly which ASes connected with which, and what the RTT was to reach either end of the cross connect.

Sanity Checks

After the data was collected, parsed, and processed, it was subjected to a series of "sanity" checks that assessed the coverage and completeness of a campaign. Relevant statistics were gathered regarding the number of IP addresses targeted, the number of traceroute measurements launched, the number of traceroute measurements that reached the targeted AS, the number that reached the targeted IP address, and the number of cross connects found.

3.4 - Inferring Cross Connect Location Using Traceroute Measurement Data

Round Trip Time (RTT) Threshold Determination

For each of the traceroute or ping measurements being sent between the vantage points and targets selected for a city, there was no guarantee that the measurement would stay in the city. However, if the RTT was low enough, it was likely that the measurement did not leave the city. This is because a low RTT can suggest that the measurement did not travel a long distance. For example, a signal that travelled for 1 millisecond is much more likely to have been sent to a destination nearby than a destination across the country. However, this threshold of what is considered low enough is different for every city. A five millisecond RTT in one city could indicate that the signal stayed within the city, whereas for a campaign in a different city it could have gone much further and left the city limits. This is because every city has different network infrastructure, size, and physical geography.

By graphing the distribution of RTTs collected by pings and traceroutes in a city, it was possible to estimate potential thresholds. Once an estimate was determined, the number was multiplied by 1.5 in order to make it more conservative.

3.5 - Validating Inferences with an External Source

PeeringDB

An important aspect of this work was using PeeringDB, a service that shows the AS tenants of a number of registered colos and IXPs. As discussed before, colos and IXPs are locations where a large number of ASes interconnect and exchange information. When a cross connect was found in the traceroute data, it was checked to see if it was corroborated by PeeringDB. This was done by checking to see if any IXPs or colos in the city hosted both of the ASes involved in the cross connect. Suppose a cross connect in a certain city involves AS X and AS Y, and using the RTT threshold, it is suspected that the cross connect is within the city. If PeeringDB indicates that AS X and AS Y are both serviced by the same IXP or coloin that city, then it makes our inference more likely.

3.6 - Data Analysis

City-Level Analysis

In order to perform the analysis of the data, three main questions were asked: does the traceroute indicate a cross connect between two ASes in this city? If so, does this cross connect fall below the RTT threshold determined for the city? Finally, does PeeringDB corroborate this evidence? Visualizations were created that showed how effective the traceroute measurements were at locating cross connects on a per-city basis. Discovered cross connects that are corroborated by PeeringDB and the RTT threshold are considered more likely to be in the particular city than a cross connect that had neither.

AS-Level Analysis

Several AS-level characterizations were produced to answer additional questions that were raised during the city-level analysis. This also demonstrates the versatility of the data set and its usefulness in coming to inferences about numerous aspects of AS interconnectivity. Visualizations were generated that focused on the connectivity of individual ASes across all cities, and was able to show how an AS interacted with other ASes, as well as the level of presence of an AS within different cities.

AS-Pair-Level Analysis

Alongside the AS-level analysis, it was also possible to observe the behavior of individual AS-pairs. With 20 ASes being investigated, there are a total of 190 possible AS-pairs in this thesis. This analysis looked at connectivity patterns between pairs of ASes, as well as the presence of cross connects between AS-pairs.

Chapter 4: Data Collection and Data Set

This chapter will describe and characterize each of the data sets that were collected. This includes data that was collected using the RIPE Atlas measurements, as well as relevant meta data such as which cities were selected, which ASes were selected, and RIPE Atlas probe information. It is important to understand that all of the data collected is not present in this paper.

4.1 - Experimental Setting

Selected Cities

City	# of Probes	Latitude	Longitude
San Francisco, California	120	37.733795	-122.446747
Washington, D.C.	84	38.889931	-77.009003
New York City, New York	74	40.730610	-73.935242
Philadelphia, Pennsylvania	72	40.071030	-75.205963
Boston, Massachusetts	60	42.364758	-71.067421
Seattle, Washington	57	47.608013	-122.335167
Rochester, New York	48	43.161030	-77.610924
Los Angeles, California	46	34.052235	-118.243683
Detroit, Michigan	31	42.331429	-83.045753
Chicago, Illinois	28	41.881832	-87.623177
Denver, Colorado	26	39.742043	-104.991531
Dallas, Texas	22	32.897480	-97.040443
Minneapolis, Minnesota	20	44.986656	-93.258133
Atlanta, Georgia	19	33.753746	-84.386330
Portland, Oregon	17	45.512794	-122.679565
Columbus, Ohio	16	40.061974	-82.991028
Raleigh, North Carolina	15	35.787743	-78.644257

Table 4.1.: Targeted cities and their RIPE Atlas probe count

Selected ASes

AS Number	# of Customer ASes	# of IPv4 Addresses	AS Name	Organization
3356	24,553	715,498,496	LEVEL3	Level 3 Communications, Inc.
174	17,891	648,411,904	COGENT-174	Cogent Communications
3257	16,963	593,284,864	TINET-BACKBONE	Tinet Spa
1299	15,743	482,770,944	TELIANET	TeliaSonera AB
2914	13,046	511,799,040	NTT-COMMUNICATIONS-2914	NTT America, Inc.
6453	9,607	447,639,040	AS6453	TATA Communications (America), Inc.
6762	8,646	238,719,232	SEABONE-NET	Telecom Italia Sparkle S.p.A.
6939	8,048	305,295,872	HURRICANE	Hurricane Electric, Inc.
2828	5,866	249,494,784	XO-AS15	XO Communications
3549	5,534	89,031,424	LVLT-3549	Level 3 Communications, Inc.
1273	4,469	147,194,624	CW	Cable and Wireless Worldwide PLC
3491	4,196	184,550,656	BTN-ASN	Beyond The Network America, Inc.
6461	3,858	106,985,984	ABOVENET	Abovenet Communications, Inc.
3320	3,527	228,420,864	DTAG	Deutsche Telekom AG
20485	3,323	28,348,672	TRANSTELECOM	Closed Joint Stock Company TransTeleCom

9002	3,271	21,806,080	RETN-AS	RETN Limited
7018	3,189	254,971,136	ATT-INTERNET4	AT&T Services Inc
701	3,058	352,277,248	UUNET	MCI Communications Services, Inc. d/b/a Verizon Business
209	2,934	231,474,688	CENTURYLINK-US-LEGACY-QWEST	Qwest Communications Company, LLC
1239	2,923	329,687,040	SPRINTLINK	Sprint

Table 4.2: Selected ASes

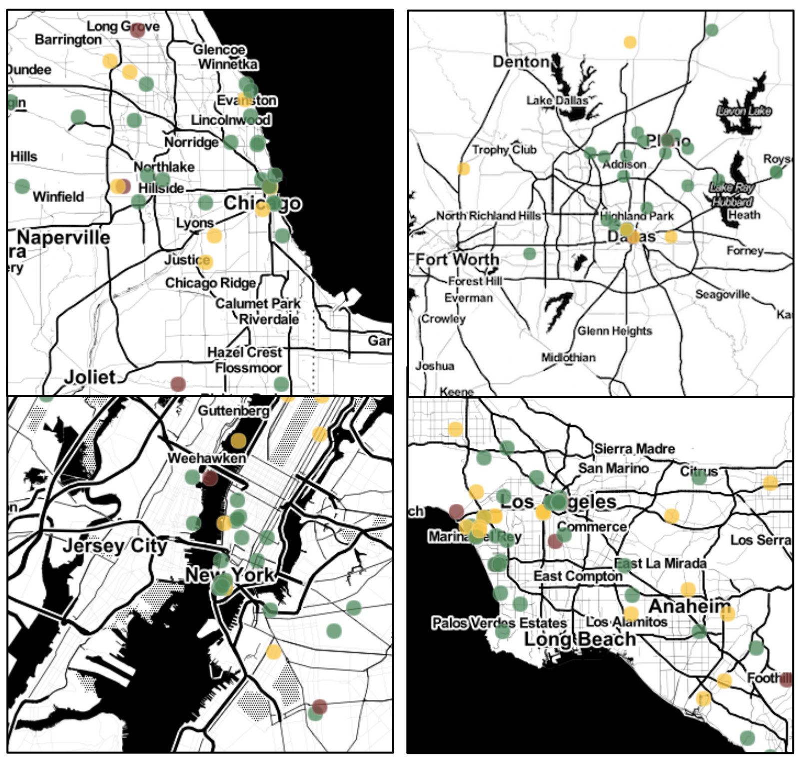

Figure 4.1: Probe coverage images (courtesy of RIPE Atlas). Each dot represents a probe. Green dots represent connected probes, yellow dots represent disconnected probes, and red dots indicate abandoned probes. Starting clockwise from the top left: Dallas, Chicago, New York, and Los Angeles

4.2 - Traceroute Dataset

The traceroute data set encompassed 17 cities and hundreds of probes. The number of probes used per city can be seen in tables 4.1. All of the traceroute data collected for this thesis was collected on May 1st, 2016.

As mentioned in the methodology, up to 3 IP addresses were selected for each of the top 20 ASes in order to increase the likelihood that a traceroute measurement succeeded. This is due to reasons that affect all traceroutes, as well as reasons that affect the RIPE Atlas platform specifically. A traceroute is launched from a vantage point to a target, but that does no ensure that it will arrive. There are numerous reasons why a traceroute might fail to reach its target, such as firewalls, non-responsive routers, or an IP address that is no longer in use. As for RIPE Atlas, there is a chance that a probe is unresponsive, disconnected, or under too much load to process a measurement request.

Because of these potential measurement failures, the campaigns for each city were examined for completeness. Tables were produced that showed whether or not the traceroutes were reaching their intended targets and returning usable results. Since showing the tables for each and every city would be superfluous in the context of this thesis, four samples are shown in figure 4.2. Each table shows 6 counts for each of the 20 targeted ASes. The first shows the number of unique prefixes targeted. This is the number of prefixes that were found inside of that AS. Up to three were selected when launching the campaigns based on availability. Notice in the tables that there are some ASes that had more than 3 prefixes selected. This was unintentional. Only three prefixes were ever selected programatically, but RIPE Atlas seems to have resolved additional addresses for some of the ASes. Since this is a count of unique prefixes, they are not duplicates and were still used for our data set.

The next value is the number of traceroutes launched, which is a multiple of the number of prefixes, since the same number of measurements were launched at each prefix. Of those launched, only a fraction will have actually reached the AS that was targeted. Furthermore, of those that reached the AS, only some will have reached the targeted IP. This gradual decrease is reflected in the tables below as well.

The last two statistics record the number of “hits” or cross connects found that were relevant to this thesis. A relevant cross connect is a cross connect between two routers, each of which is registered to one of the top 20 ASes. Targeted hits indicate a cross connect between two top 20 ASes where one of the ASes was the AS towards which the traceroute was launched. Opportunistic hits are cross connects between top 20 ASes that were found during traceroute measurements that were not launched towards either of those ASes.

Traceroute Breakdown per AS for the city of NewYorkCity

	3356	174	3257	1299	2914	6453	6762	6939	2828	3549	1273	3491	6461	3320	20485	9002	7018	701	209	1239	Total
Unique Prefixes Targeted	6	3	3	3	3	3	3	3	3	0	3	3	3	3	0	0	3	3	3	3	54
Traceroutes Launched	444	222	222	222	222	222	222	222	222	0	222	222	222	222	0	0	222	222	222	222	3996
Traceroutes Reached AS	343	197	146	219	219	218	219	216	186	0	219	218	187	219	0	0	219	184	132	0	3341
Traceroutes Reached IP	219	177	146	219	219	146	218	216	73	0	219	218	146	219	0	0	0	73	0	0	2508
Targeted Hits	39	28	34	50	36	41	49	52	52	0	54	54	27	54	0	0	48	47	51	0	716
Opportunistic Hits	87	19	45	48	30	33	42	45	45	0	48	45	21	45	0	0	45	6	45	45	694

Traceroute Breakdown per AS for the city of LosAngeles

	3356	174	3257	1299	2914	6453	6762	6939	2828	3549	1273	3491	6461	3320	20485	9002	7018	701	209	1239	Total
Unique Prefixes Targeted	6	3	3	3	3	3	3	3	3	0	3	3	3	3	0	0	3	1	3	3	52
Traceroutes Launched	300	150	150	150	150	150	150	150	150	0	150	150	150	150	0	0	150	50	150	150	2600
Traceroutes Reached AS	217	147	148	99	147	107	150	142	105	0	150	150	149	88	0	0	107	48	142	54	2150
Traceroutes Reached IP	149	50	148	99	147	96	150	99	49	0	150	150	149	88	0	0	50	48	99	54	1775
Targeted Hits	14	18	18	12	9	11	18	10	10	0	18	15	18	11	0	0	16	1	17	4	220
Opportunistic Hits	22	13	13	17	13	17	12	12	14	0	15	11	13	13	0	0	9	2	13	16	225

Traceroute Breakdown per AS for the city of Dallas

	3356	174	3257	1299	2914	6453	6762	6939	2828	3549	1273	3491	6461	3320	20485	9002	7018	701	209	1239	Total
Unique Prefixes Targeted	6	3	3	3	2	3	1	3	3	0	0	1	3	2	0	0	3	0	3	3	42
Traceroutes Launched	132	66	66	66	44	66	22	66	66	0	0	22	66	44	0	0	66	0	66	66	924
Traceroutes Reached AS	63	63	63	42	42	57	3	63	45	0	0	21	63	7	0	0	63	0	43	43	681
Traceroutes Reached IP	63	42	63	42	42	42	0	42	0	0	0	21	21	7	0	0	4	0	21	43	453
Targeted Hits	7	24	24	16	12	17	2	24	14	0	0	7	15	0	0	0	18	0	7	22	209
Opportunistic Hits	19	10	9	12	6	12	3	15	9	0	0	3	10	7	0	0	3	0	9	13	140

Traceroute Breakdown per AS for the city of Chicago

	3356	174	3257	1299	2914	6453	6762	6939	2828	3549	1273	3491	6461	3320	20485	9002	7018	701	209	1239	Total
Unique Prefixes Targeted	6	3	3	3	3	3	2	3	3	0	0	3	2	3	0	0	3	3	2	3	48
Traceroutes Launched	150	75	75	75	75	75	50	75	75	0	0	75	50	75	0	0	75	75	50	75	1200
Traceroutes Reached AS	109	75	72	75	75	75	50	44	75	0	0	75	50	33	0	0	29	12	50	0	899
Traceroutes Reached IP	100	75	72	75	75	75	50	33	75	0	0	75	50	33	0	0	28	0	50	0	866
Targeted Hits	2	6	6	6	6	6	4	6	6	0	0	6	4	4	0	0	0	0	2	0	64
Opportunistic Hits	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Figure 4.2: Tables of traceroute validation statistics for Chicago, Dallas, Los Angeles, and New York

Table 4.3 shows the number of traceroutes launched in each city, and the total number of traceroutes launched for this thesis. This gives a sense of the total size of the data set.

City	Traceroutes Launched
Seattle, Washington	3036
Denver, Colorado	1134
San Francisco, California	5740
Los Angeles, California	2600
Portland, Oregon	544
Boston, Massachusetts	60
Dallas, Texas	924
Raleigh, North Carolina	280
Chicago, Illinois	1200
New York City, New York	3996
Washington, D.C.	4644
Minneapolis, Minnesota	641
Columbus, Ohio	395
Atlanta, Georgia	828
Rochester, New York	928
Philadelphia, Pennsylvania	3942
Detroit, Michigan	928
Total	**34769**

Table 4.3: Number of traceroutes launched for each city

Once the traceroutes were collected and processed, it was possible to view all of the cross connects and filter by whether or not they were a cross connect that involved the top 20 ASes. By looking at the traceroute log, it was possible to see the RTT of the far end of the cross connect, which is the RTT in particular that was used to find the RTT threshold.

The data used for this thesis was very minimally cleaned. It was determined that an RTT over 100 milliseconds was not related to the distance that the signal had to travel, but rather indicative of other issues such as queueing delays or network congestion that were unimportant for the purposes of this thesis. This eliminated some of the outliers and made the RTT data fit within a range of 0 to 100 milliseconds.

4.3 - Ping Dataset

Table 4.4 shows the number of pings launched for each city. The largest number of pings possible and scheduled for launch was the N(N-1)/2, where N is the number of probes in a city. Ping measurements were also filtered, with pings with RTTs over 100 milliseconds being removed.

City	**Pings Launched**
Seattle, Washington	1514
Denver, Colorado	351
San Francisco, California	6403
Los Angeles, California	1225
Portland, Oregon	136
Boston, Massachusetts	1691
Dallas, Texas	231
Raleigh, North Carolina	91
Chicago, Illinois	304
New York City, New York	2703
Washington, D.C.	3655
Minneapolis, Minnesota	208
Columbus, Ohio	115
Atlanta, Georgia	153
Rochester, New York	955
Philadelphia, Pennsylvania	2629
Detroit, Michigan	423
Total	**22787**

Table 4.4: Number of pings launched for each city

4.4 - PeeringDB Data Set

PeeringDB is a resource that can be crawled using an API. The entire data set is capable of being downloaded into an easily parsed JSON file. Each of the colos and IXPs within each city were identified and then searched to see if they had one of the top 20 ASes as a tenant. If two top 20 ASes were found in the same building, it was noted as a potential relationship that could be detected by a traceroute measurement.

PeeringDB is not a final authority on the location of these facilities, or the presence of their tenants. Much of the data is contributed voluntarily. Some ASes, such as AS3356 and AS174, do not report their presence to PeeringDB. As such, PeeringDB cannot be treated as a complete data set. However, when it does report the presence of an AS, that information was trusted. The PeeringDB data used for this thesis was current as of May 1st, 2016.

4.5 - RTT Threshold Data Set

The type of chart used for determining the RTT threshold is a cumulative distribution function (CDF) of the RTT values of the ping and traceroute measurements. This chart tells the story about how often a certain RTT value was seen in the data. Any time the chart flattens, that means that there aren't many RTTs of that value. The idea is that communication within the city and communication outside the city should be separated by a flattening, or a 'knee', in the graph. Sections where the graph levels off can be an indicator of this threshold, with RTT values before the 'knee' being within the city, and RTT values after, being outside of the city. In this paper, we show the RTT distributions for the four cities that we have been following so far. A brief explanation will be provided to justify the selection of the RTT threshold for that city.

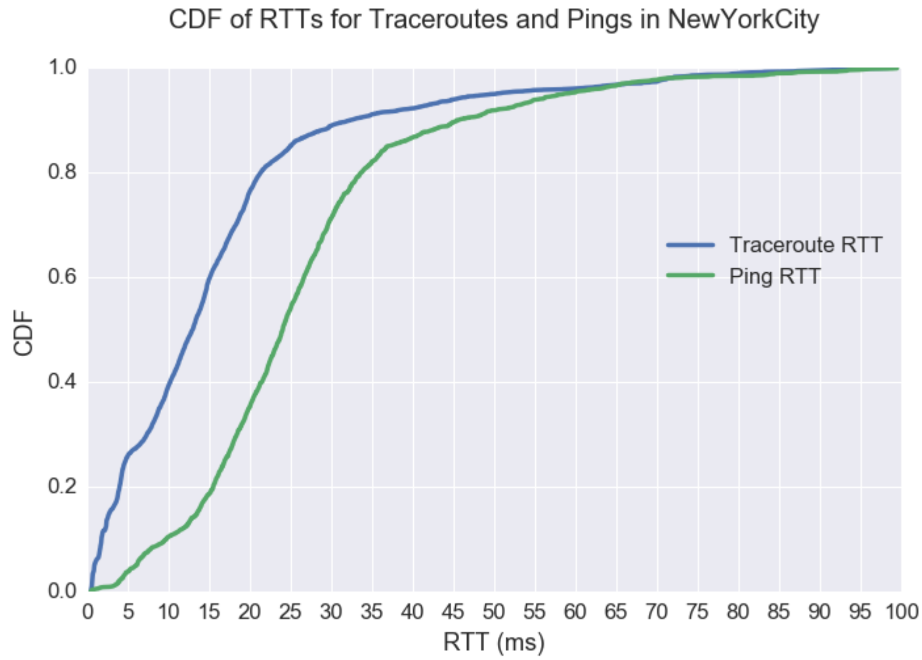

Figure 4.3: RTT CDF of both traceroute and ping measurements in New York City

New York has a slight knee at approximately the same point on both the ping RTT line and the traceroute RTT line. Thus we selected 6 milliseconds as the point where the threshold occurs and multiplied by 1.5 to achieve the final threshold of 9 milliseconds.

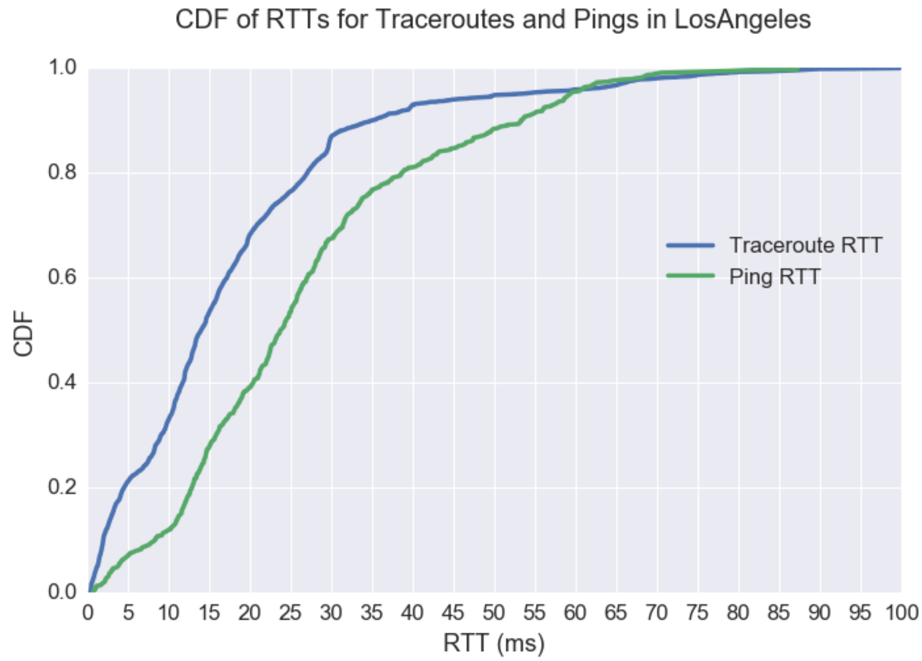

Figure 4.4: RTT CDF of both traceroute and ping measurements in Los Angeles

Los Angeles has a distribution similar to that of New York. It also exhibits a flatter slope at approximately 6 milliseconds. Therefore, the RTT threshold was set to the same as New York City's, at 9 milliseconds.

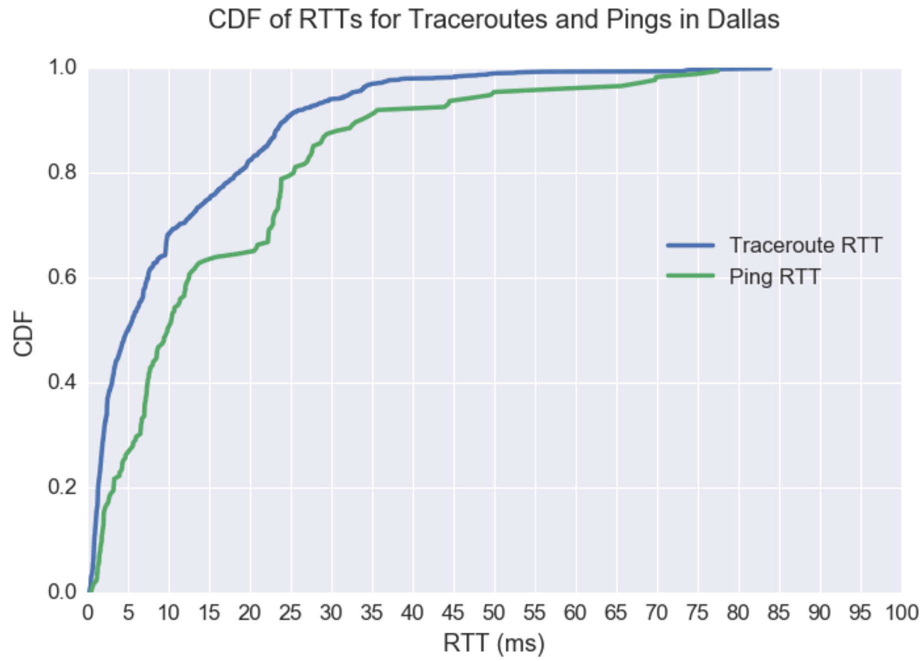

Figure 4.5: RTT CDF of both traceroute and ping measurements in Dallas

The distribution of RTT in Dallas is particularly poor for our purposes when compared to the other cities. There is no clear flattened part except for on the ping distribution, where it flattens at about 13 milliseconds. Dallas is an example of a city that did not respond well to our methodology, as the resulting RTT threshold of 19.5 milliseconds is higher relative to the threshold for every other city. Almost every cross connect found during the Dallas traceroute campaigns was expected to have an RTT below 19.5 milliseconds.

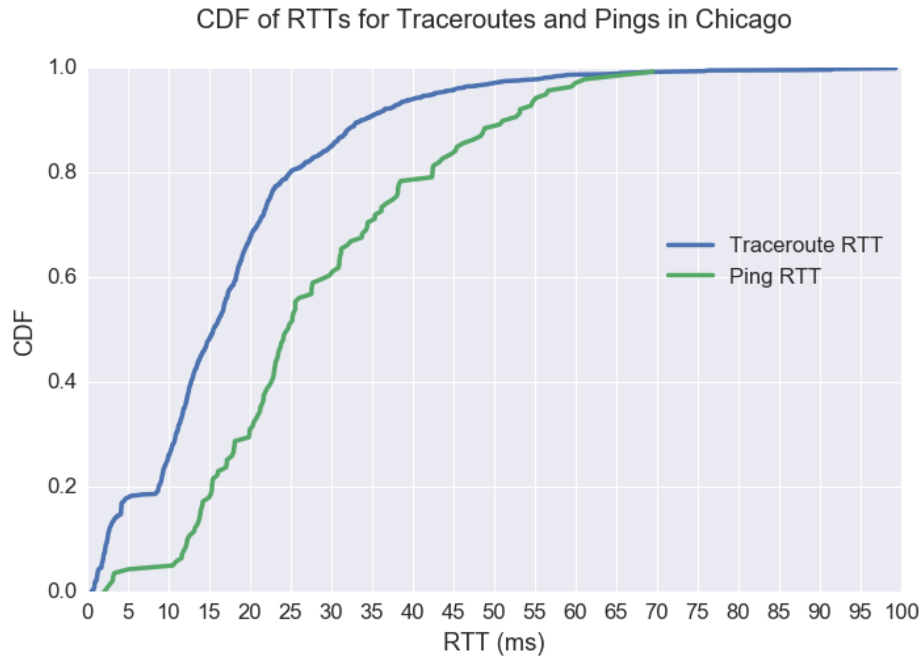

Figure 4.6: RTT CDF of both traceroute and ping measurements in Dallas

In contrast to Dallas, Chicago is a city that responded extremely well to our method. The ping distribution and traceroute distribution correspond very closely, and there is a portion of the graph that is clearly flat. This RTT threshold is very clear and creates a divide between two ranges of RTT. Table 4.5 on the next page shows the RTT threshold estimates for every city in this study.

City	RTT Threshold Estimate (ms)
Seattle, Washington	6
Denver, Colorado	9
San Francisco, California	12
Los Angeles, California	9
Portland, Oregon	9
Boston, Massachusetts	8.5
Dallas, Texas	19.5
Raleigh, North Carolina	9
Chicago, Illinois	7.5
New York City, New York	9
Washington, D.C.	10.5
Minneapolis, Minnesota	10
Columbus, Ohio	15
Atlanta, Georgia	7.5
Rochester, New York	10
Philadelphia, Pennsylvania	7.5
Detroit, Michigan	9

Table 4.5: Estimated RTT threshold for each city

Chapter 5: City-Level Analysis

The following chapter will consist of analyses of the same four cities that we have been using as examples in earlier chapters: Los Angeles, Dallas, Chicago, and New York City. What this section aims to provide is an understanding of how the data collected can be used and visualized.

There are several main motivating questions for the city-level analysis.

1. How many cross connects found are evidenced to be in the city by their RTT? How many of these are corroborated by PeeringDB?
2. Of the pairs of ASes that PeeringDB indicates are present in the city, how many have established a cross connect?
3. Which ASes of the top 20 ASes has the largest presence in each city? Of those ASes, which are generally serving as a facilitator (found at the front end of a cross connect) and which generally serve as a recipient (found at the back end of a cross connect).

The following sections will answer these questions for the four cities selected by utilizing diagrams and tables derived from the data. While these diagrams are presented only for these four cities, each city had these diagrams produced and they are available in the supplementary materials.

5.1 - Number of Cross Connects Found & Evidence of AS Presence

The first diagram is a matrix that shows the relationships that were found within the city. This matrix does not show the number of cross connects found, but rather that at least one was found, thus indicating the presence of a relationship. If an entire cell is filled with grey, that means that a cross connect was found. If the box has a red marker in it, then this means that at least one of the discovered cross connects had an RTT below the RTT threshold for that city. This suggests that the relationship between the ASes exists within the city. Finally, a green marker indicates that the PeeringDB record shows that both ASes are present in a colo or IXP within the city. Note that this marker can be placed whether a cross connect was found or not: hence light green markers show that there is only PeeringDB evidence for both ASes being within the city. A dark green marker indicates that it has been placed over a grey marker, indicating the presence of a cross connect. Therefore, any cell that has both a red marker and a dark green marker is indicative of a high level of confidence that the ASes in question do have a cross connect within the city.

Note that the matrix is not symmetric. The y-axis indicates the AS that was first encountered in a cross connect. If there is a cross connect from A to B, the y-axis shows A. The x-axis shows the AS that was encountered at the other end of the cross connect, or B.

We provide a table that shows the number of AS-pairs that had established a cross connect. It also shows what proportion of these AS-pairs had RTT based evidence, PeeringDB evidence, or both forms of evidence to place them within the city.

Los Angeles, CA, US

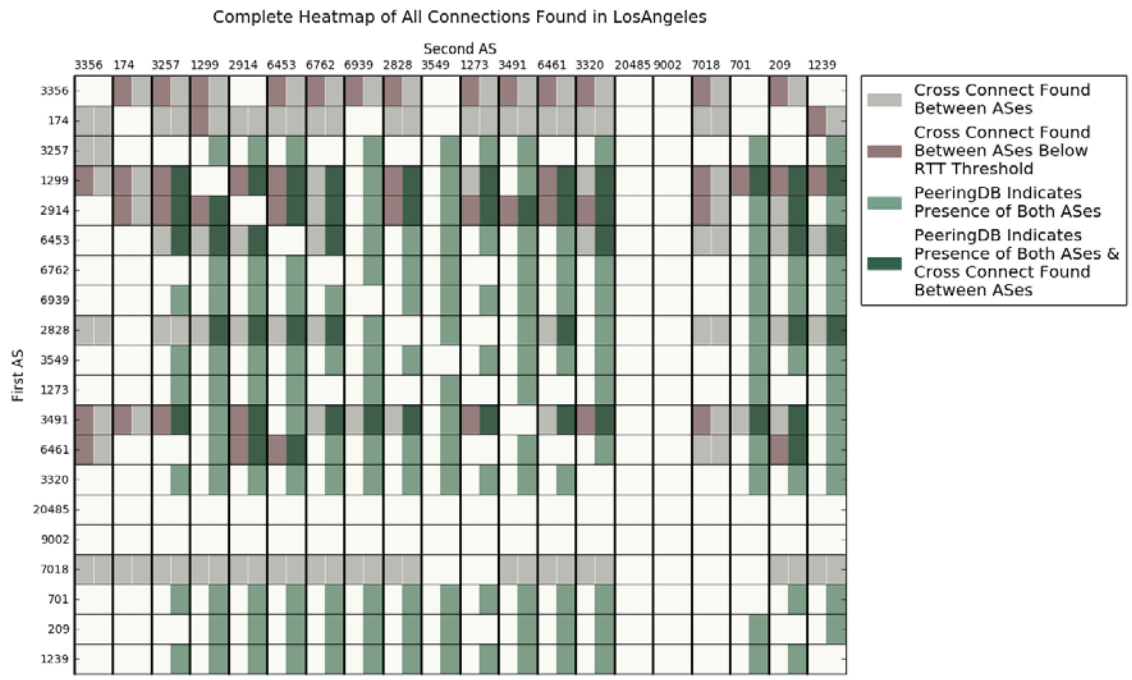

Figure 5.1: Matrix of detection and corroboration of cross connects in Los Angeles

Amount of AS-pairs With A Cross Connect	**Amount with No Additional Evidence**	**Amount with Only RTT Evidence**	**Amount with Only PeeringDB Evidence**	**Amount with Both Forms of Evidence**
103	28 (27.18%)	27 (26.21%)	22 (21.36%)	26 (25.24%)

Table 5.1: Cross connect evidence summary for Los Angeles

As can be seen in figure 5.1, Los Angeles has a large density of colos and IXPs, with 15 of the top 20 ASes having a presence in the city. Table 5.1 formats this information into numbers: of the 103 AS-pairs that had a cross connect, 26.21% had just RTT evidence, and 25.24% had both RTT evidence and PeeringDB evidence of their presence within the city. This is over 50% with RTT evidence, which splits the total number of relationships with a cross connect approximately in half. All in all, over 70% of relationships had some form of evidence that they were present within the city. As can be clearly seen in the top row of figure 5.1, AS3356 in particular has all of its relationships corroborated by RTT evidence, which might suggest that it has a large presence within Los Angeles.

Dallas, TX, US

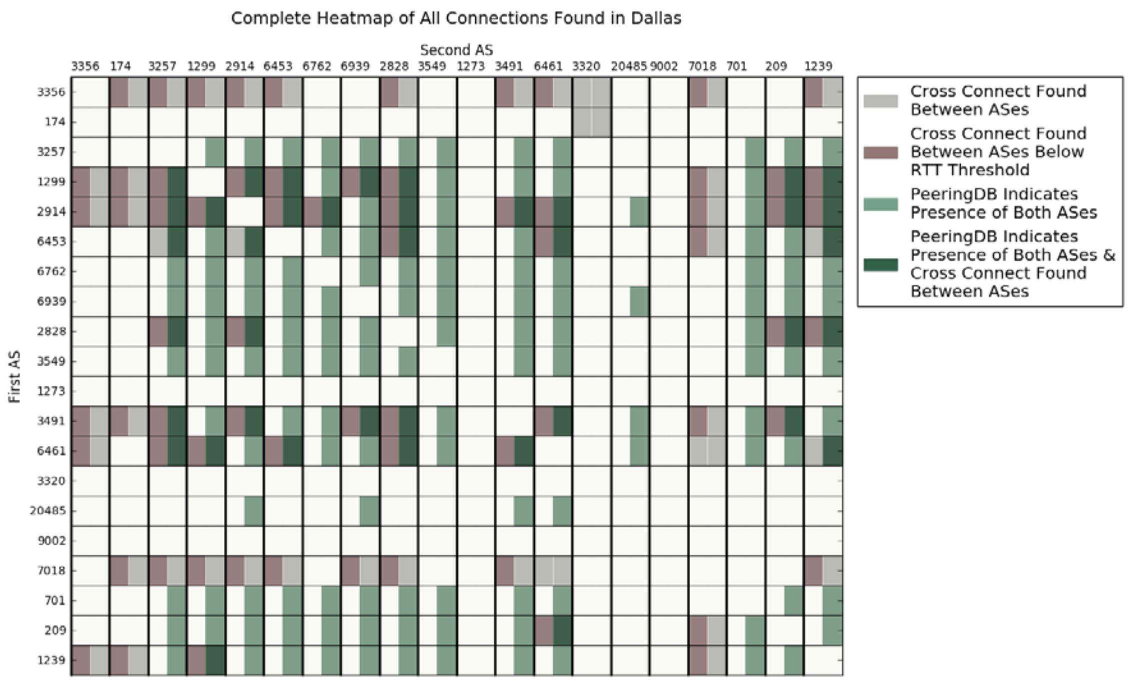

Figure 5.2: Matrix of detection and corroboration of cross connects in Dallas

Amount of AS-pairs With A Cross Connect	Amount with No Additional Evidence	Amount with Only RTT Evidence	Amount with Only PeeringDB Evidence	Amount with Both Forms of Evidence
77	4 (5.19%)	34 (44.16%)	4 (5.19%)	35 (45.45%)

Table 5.2: Cross connect evidence summary for Dallas

In Dallas, the story is quite different than in Los Angeles. If you recall, Dallas has a relatively high RTT threshold in comparison to other cities. This example serves mostly to show how a very high RTT threshold can skew the data. Dallas also has a very large number of top 20 ASes found together in its colos and IXPs. Table 5.2, however, indicates that PeeringDB might be the only reliable evidence here. Of the 77 AS-pairs that established a cross connect, nearly 90% fell below the RTT threshold. Given the somewhat random nature of traceroutes, it is unlikely that nearly all of the measurements done for Dallas remained within the city. It is difficult to make any definitive statements or observations about this data regarding the RTT. However, the PeeringDB evidence indicates that there is a strong colo and IXP presence in Dallas.

Chicago, IL, US

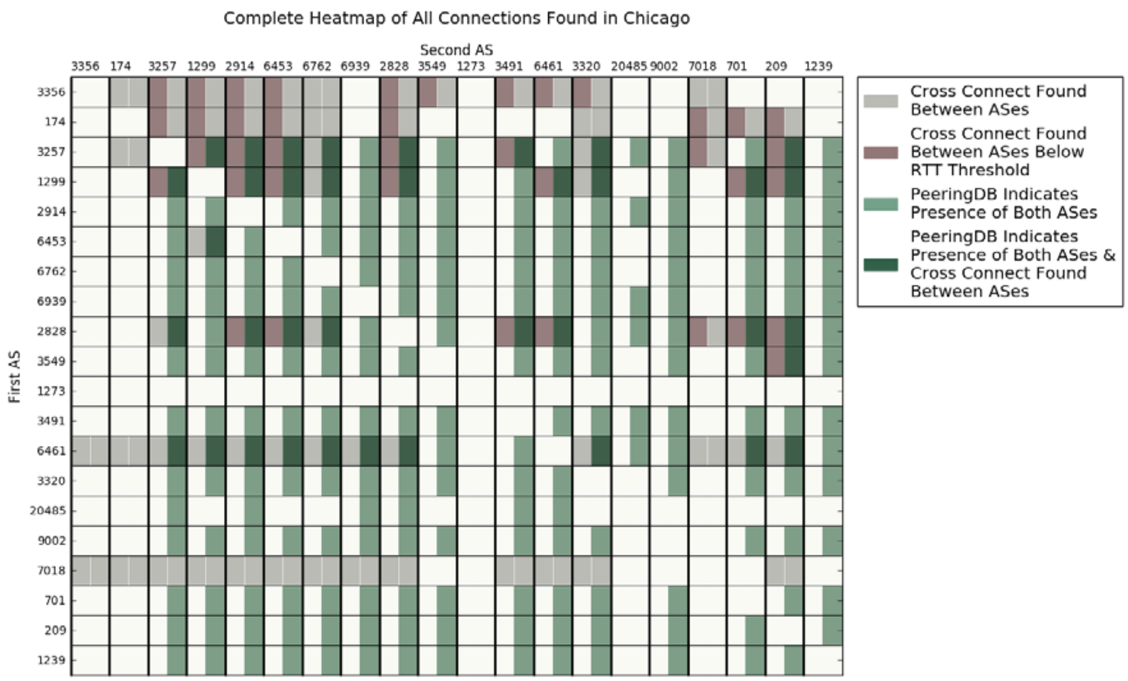

Amount of AS-pairs With A Cross Connect	Amount with No Additional Evidence	Amount with Only RTT Evidence	Amount with Only PeeringDB Evidence	Amount with Both Forms of Evidence
78	22 (28.21%)	19 (24.36%)	17 (21.79%)	20 (25.64%)

Figure 5.3: Matrix of detection and corroboration of cross connects in Chicago

Table 5.3: Cross connect evidence summary for Chicago

Chicago has a threshold that inspires quite a bit of confidence: after all, as seen in figure 4.6, the RTT distribution for Chicago very clearly indicated an RTT threshold according to out methodology. The top 20 ASes have a strong presence in colos and IXPs, with 17 ASes reporting their presence according to PeeringDB. This is more than Los Angeles and Dallas. Chicago even has PeeringDB data regarding two ASes that were absent in Los Angeles: AS20485 and AS9002. While no cross connects were found for those ASes, it indicates that Chicago is a hub of top 20 AS activity.

Table 5.3 has a very similar profile as Los Angeles. Approximately 70% of AS relationships that had established a cross connect had evidence placing them within the city. Once more, about 50% were corroborated by being below the RTT threshold. The fact that Chicago had such a clear RTT threshold and has such similar percentages as Los Angeles suggests a pattern for a well selected RTT threshold. As in Los Angeles AS3356 has many of its relationships corroborated by RTT evidence. AS3356, alongside AS174, AS3257, and AS1299, all seem to have a strong presence within the city.

New York City, NY, US

Amount of AS-pairs With A Cross Connect	Amount with No Additional Evidence	Amount with Only RTT Evidence	Amount with Only PeeringDB Evidence	Amount with Both Forms of Evidence
127	28 (22.05%)	47 (37.01%)	12 (9.45%)	40 (31.50%)

Figure 5.4: Matrix of detection and corroboration of cross connects in New York City

Table 5.4: Cross connect evidence summary for New York City

New York City is similar to Los Angeles in terms of both size and connectivity. It is also home to quite a few colos, with 192 pairs of ASes being present in New York in colos or IXPs according to PeeringDB. It exhibited a similar RTT threshold as well.

127 cross connects were found using traceroutes. However, unlike Chicago and Los Angeles, which both had about 50% of cross connects fall below their respective

RTT thresholds, New York City has over 68% that fall below the threshold. This could be indicative of New York City's importance for connectivity in the region.

Summary

- Between 70% and 95% (with Dallas serving as an outlier) of relationships with a cross connect had corroborating evidence that the relationship exists within the city.
- Chicago and Los Angeles had 50% of its relationships exhibit a cross connect that fell below the RTT threshold. New York City had 68% fall below, and Dallas had almost 90%. In the case of New York, it could be suggestive of more AS-pairs choosing to establish cross connects there. In the case of Dallas, it is most likely a symptom of an RTT threshold that is too high.
- All four of these cities are extremely well connected. In the case of Dallas, however, the RTT threshold is so high that it casts doubt on the validity of the observations made for that city.

5.2 - How Many AS-Pairs Present in a City According to Peering DB Have Established a Cross Connect?

In this section we explore the rate at which top 20 ASes take advantage of opportunities to interconnect with other top 20 ASes in a city. These opportunities are defined as any moment PeeringDB indicates that both ASes are present in the same colo or IXP within a city. Once the ASes are in the same building, they may or may not decide to establish a cross connect based on a variety of factors. Table 5.5 shows the number of AS-pairs found within a city by PeeringDB. It then provides the amount of those that established a cross connect, and the amount that established a cross connect below the RTT threshold.

City	Amount of AS-pairs Found by PeeringDB	Amount of Cross Connects That Occurred Where PeeringDB Found AS-pair	Amount of Cross Connects That Occurred Where PeeringDB Found AS-pair Below The RTT Threshold
Los Angeles, CA, US	**97**	48 (49.48%)	23 (23.71%)
Dallas, TX, US	**82**	39 (47.56%)	35 (42.68%)
Chicago, IL, US	**111**	37 (33.33%)	20 (18.02%)
New York City, NY, US	**96**	52 (54.17%)	40 (41.67%)

Table 5.5: PeeringDB evidence summary for Los Angeles, Dallas, Chicago, and New York

Of the 190 possible AS-pairs (not accounting for direction), PeeringDB indicated that Los Angeles contained 97, or about 51% of the total possible. Of those, only 49.48% had any sort of indication that a cross connect had been established, and only 23.71% had a cross connect that fell below the RTT threshold. This suggests that while there is the potential for cross connects between certain ASes to be established in the colos and IXPs of Los Angeles, these ASes might have chosen to not create them. This might be because ASes utilize the colos and IXPs in Los Angeles to distribute data to lower tier ASes, rather than peer with other tier 1 ASes. It will be interesting to see if this behavior is exhibited in other cities.

Dallas exhibited 82 of the possible 190 AS-pairs, so 43%. Much like Los Angeles, 47.56% of these pairs had a cross connect found, but a much larger proportion of 42.68% fell below the RTT threshold. This further indicates that Dallas had an RTT threshold that was most likely too high. It is difficult to draw any conclusions from this data.

Of the 111 AS-pairs that PeeringDB said were in Chicago, 18.02% had cross connects between them that fell below Chicago's RTT threshold. This is much more similar to Los Angeles' value (23.71%). It once again suggests that top 20 ASes aren't utilizing colos and IXPs to connect with other top 20 ASes.

New York City so far seems to be more similar to Dallas in terms of the validity of its RTT threshold. Like Dallas, there were very few AS relationships that had no evidence. In table 5.5 above, it can be seen that once again New York City is very similar to Dallas in that it also has a much higher rate of AS-pairs in colos and IXPs having an established cross connect. This could be attributed to one of two things: either the RTT threshold was too generous for New York (such as is most likely the case with Dallas), or indeed more top 20 ASes choose to take advantage of the opportunity provided by IXPs and colocation facilities to connect with other top 20 ASes in New York City. This might be the case, as top 20 ASes might not need to connect to other top tier ASes in multiple cities. New York City's status as one of the largest and most interconnected cities in the world could make it an ideal candidate for establishing a cross connect with another top 20 AS.

Summary

- In some cities, such as Chicago and Los Angeles, top 20 ASes only take advantage of approximately 20% of potential cross connects within colos and IXPs.

- While Dallas seems to have an RTT threshold that is too high, the story might be different in New York City, despite the similar numbers. It could be the case that, due to the importance and size of New York City, top 20 ASes choose to utilize more of the potential connections within colos and IXPs.

5.3 - Which ASes Have the Largest Presence in a City?

In this section, we look at the AS presence within a single city, and see if there is any particular AS or AS relationship that dominates the landscape of a particular city. These are large ASes, many of them competing with each other for customers and coverage. Understanding what ASes are present in a city, and to what extent, is important.

To that end, we will be taking a look at two different charts. The first is a bar graph that indicates the number of cross connects found in each city. The y-axis indicates the count, and the x-axis indicates the AS in question. The bars are color coded, with blue indicating a cross connect sending data from the AS, and green indicating a cross connect sending data to the AS. Any AS that has more connections going from it can be considered a facilitator, in that it sends data to other ASes. An AS with a mostly green bar indicates that it is a recipient of data from another AS. The second chart is a CDF diagram, this time showing the number of cross connects found for each relationship, of which there are 190 possible. The x-axis indicates the number of cross connects found. From this chart it is possible to see how many AS-pairs did not have any cross connects in the city.

Los Angeles, CA, US

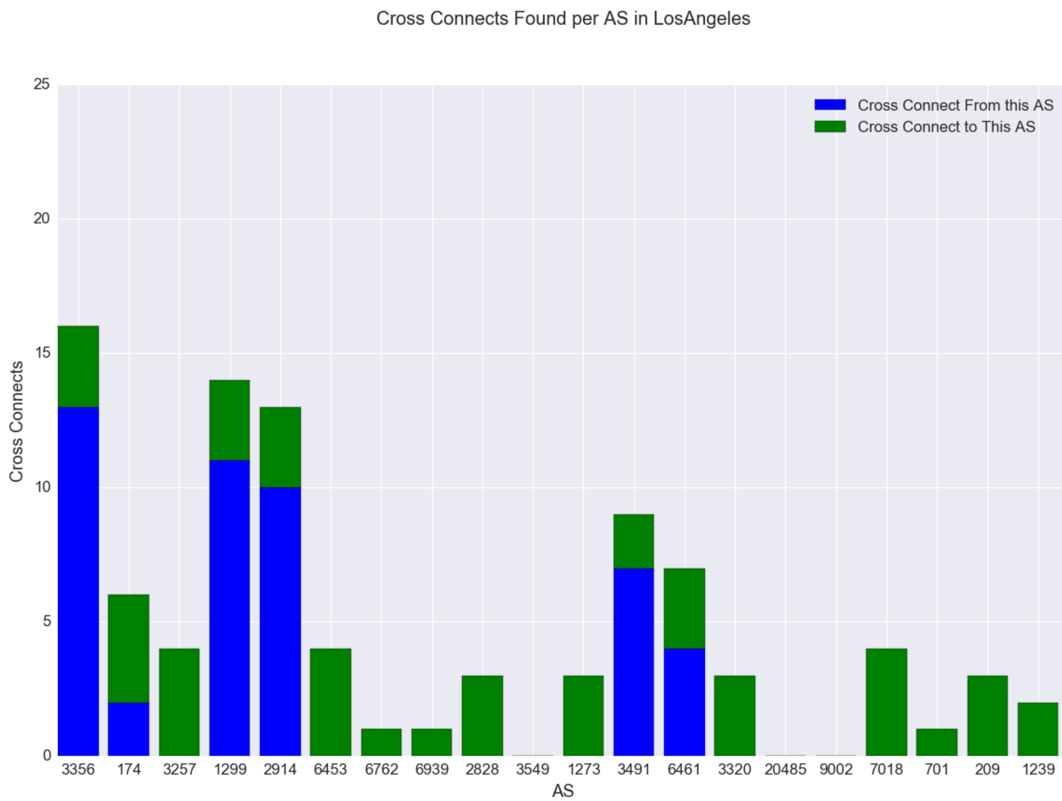

Figure 5.5: Bar chart of AS involvement in cross connects in Los Angeles

From figure 5.5, it can be seen that several ASes have quite a few more cross connects than others. They are not spread out evenly. In particular, AS3356, AS1299, and AS2914 were involved in the most cross connects. Furthermore, they can be certainly considered as facilitators since they are predominantly blue. This means that for most of the cross connects in which they are involved, they are sending the data. Apart from a few other ASes, the rest of the ASes seem to mostly be involved as recipients.

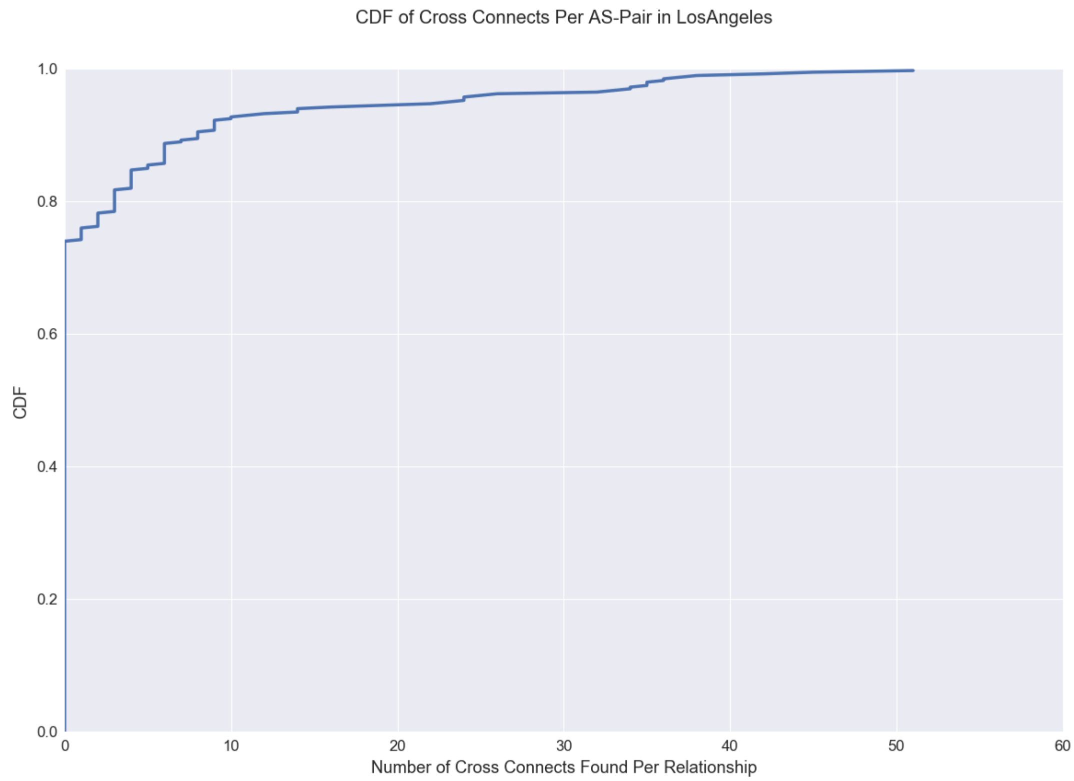

Figure 5.6: CDF of cross connects found for each AS-pair in Los Angeles

Now, as can be seen in figure 5.6, this trend continues when looking at how cross connects are distributed amongst AS-pairs. Nearly 80% of potential AS-pairs have no manifestation in Los Angeles, and approximately 20% of all AS-pairs have between 1 and 10 cross connects found. However, it also seems that there is a top 5% or so of AS-pairs that have significantly more cross connects than others. Approximately 5% of all AS-pairs range from 20 to 50 cross connects per pair. Based on the bar chart, its most likely that AS-pairs involving AS3356, AS1299, and AS2914 make up this 5%.

Dallas, TX, US

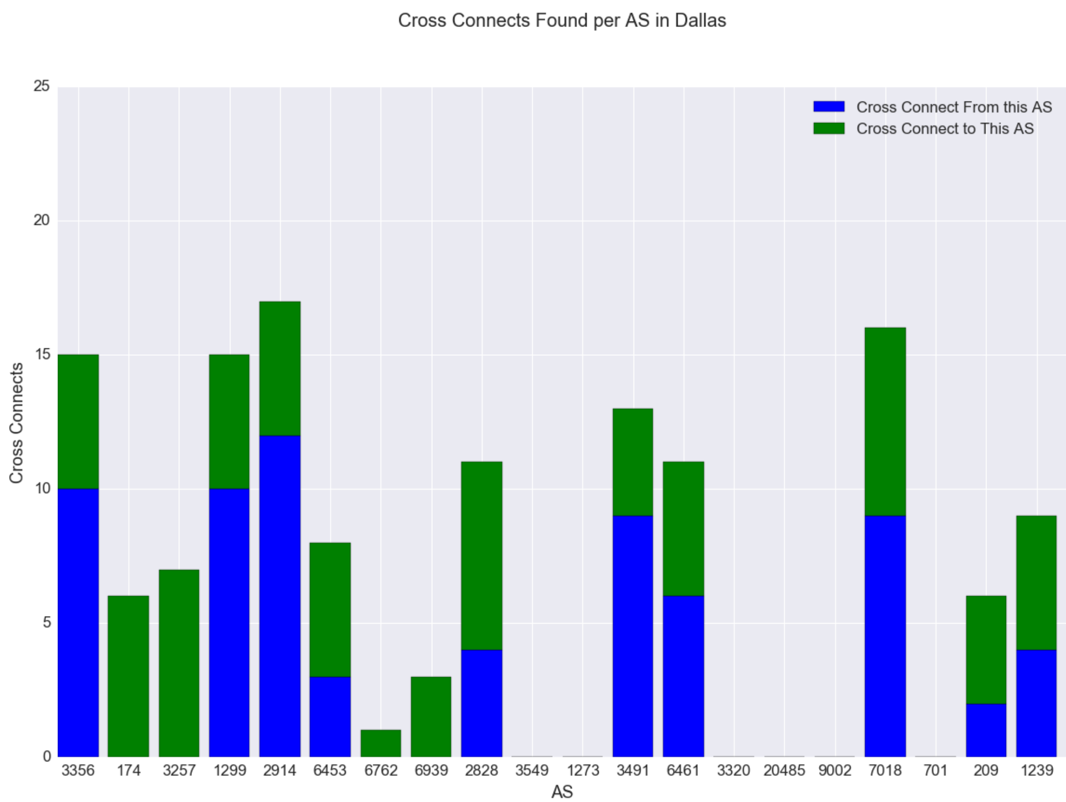

Figure 5.7: Bar chart of AS involvement in cross connects in Dallas

In Dallas, there are a number of ASes that can be viewed as facilitators and have a large number of cross connects. Once again, AS3356, AS3257, and AS1299 stand out, as well as AS7018 and AS3491. The rest of the ASes generally seem to serve as recipients. Once again, due to the fact that the vast majority of cross connects found during the Dallas campaign cleared the RTT threshold, it is not necessarily certain that these cross connects are even occurring within the city.

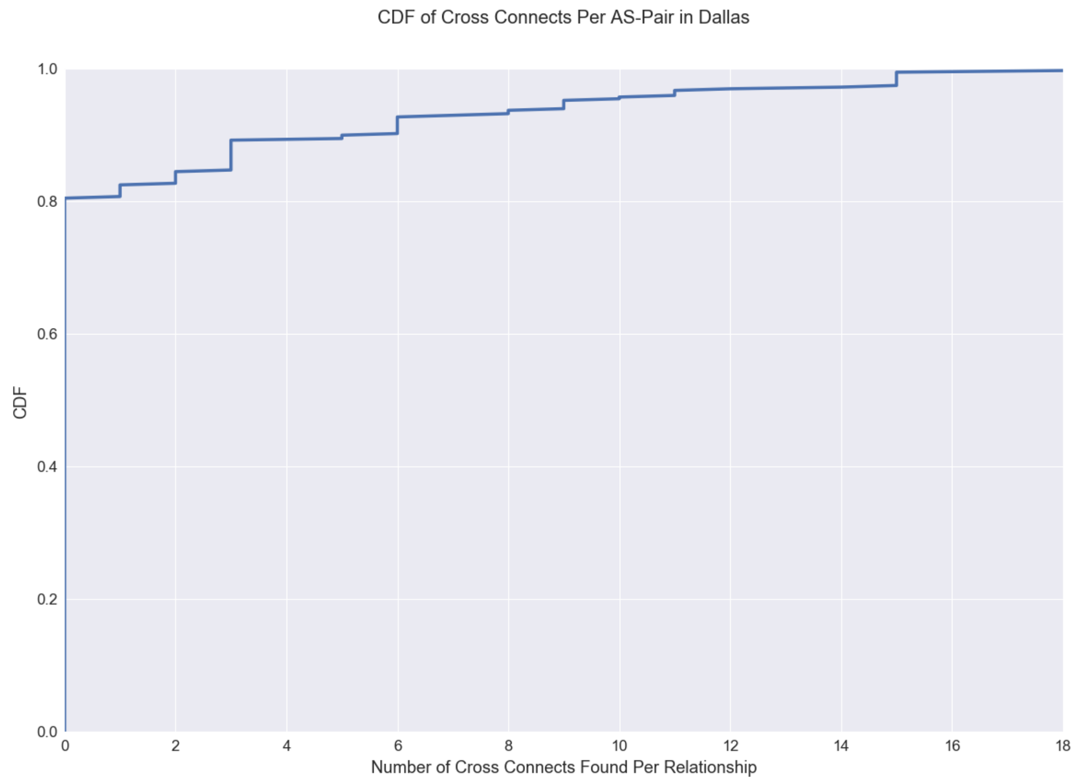

Figure 5.8: CDF of cross connects found for each AS-pair in Dallas

Figure 5.6 shows that about 80% of possible AS-pairs did not have a cross connect. About 10% of pairs have between 1 and 6 cross connects, and the next 10% have between 7 and 18. This is less of a difference than Los Angeles, and suggests that there aren't AS-pairs that exhibit disproportionately more cross connects in Dallas.

Chicago, IL, US

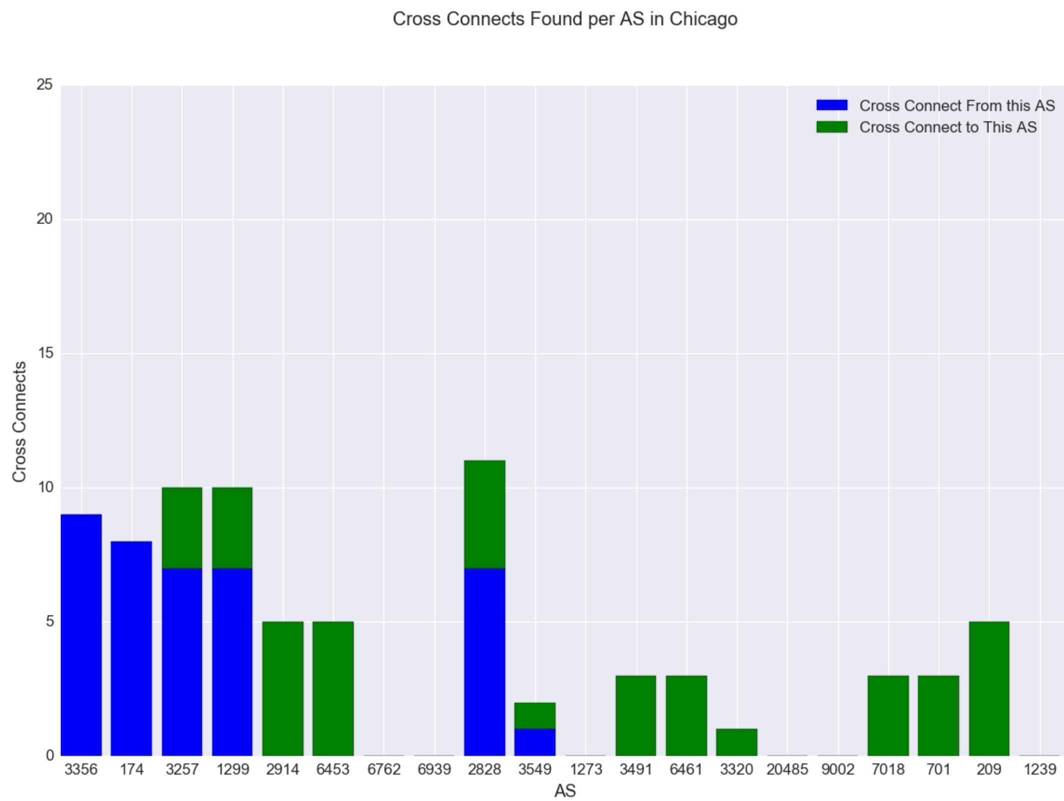

Figure 5.9: Bar chart of AS involvement in cross connects in Chicago

Chicago doesn’t have any AS with a disproportionate amount of cross connects. Rather, the ASes are split into two different groups. The first group, which includes AS3356, AS174, AS3257, AS1299, and AS2828, all have close to 10 cross connects and are all facilitators. The second group which includes the rest of the present ASes, remains around 5 cross connects and are all recipients.

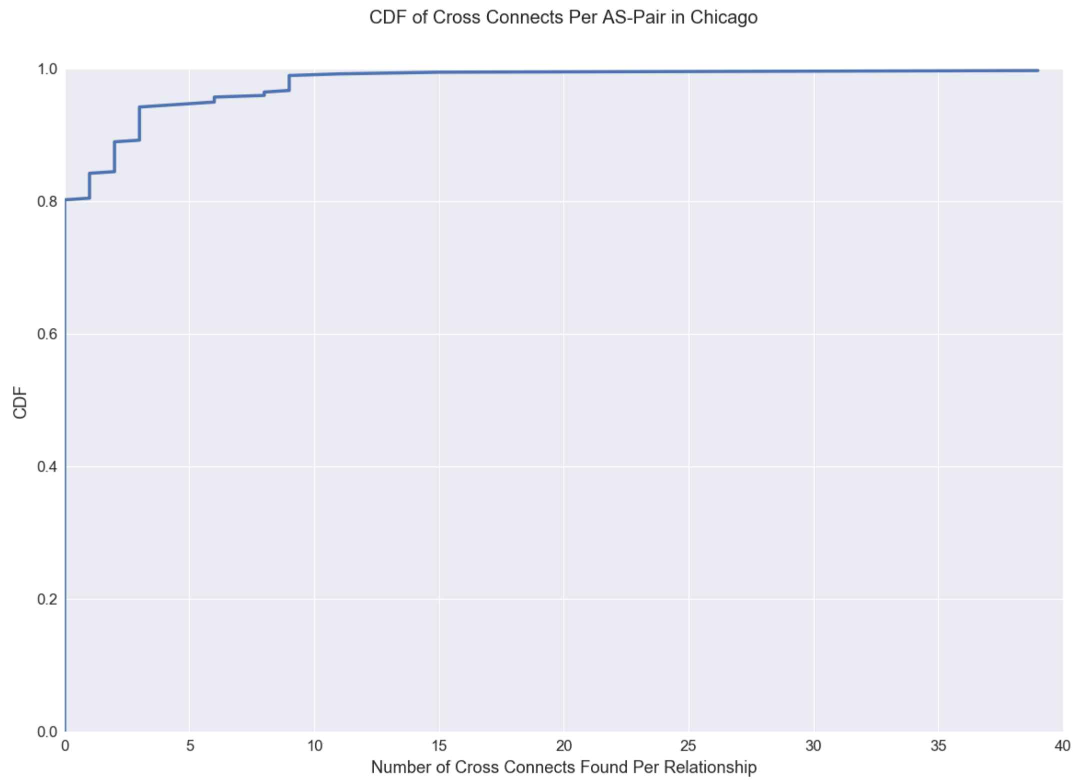

Figure 5.10: CDF of cross connects found for each AS-pair in Chicago

This grouping is also visible in figure 5.10. Once again, about 80% of AS-pairs exhibit no cross connects. The divide between the ASes with five or fewer cross connects and the ASes with more than five cross connects is clearly visible where the graph flattens out. About 15% of the AS-pairs have between 1 and 5 cross connects. About 5% have more than that, and are presumably connections involving the first group of ASes described in the previous paragraph.

New York City, NY, US

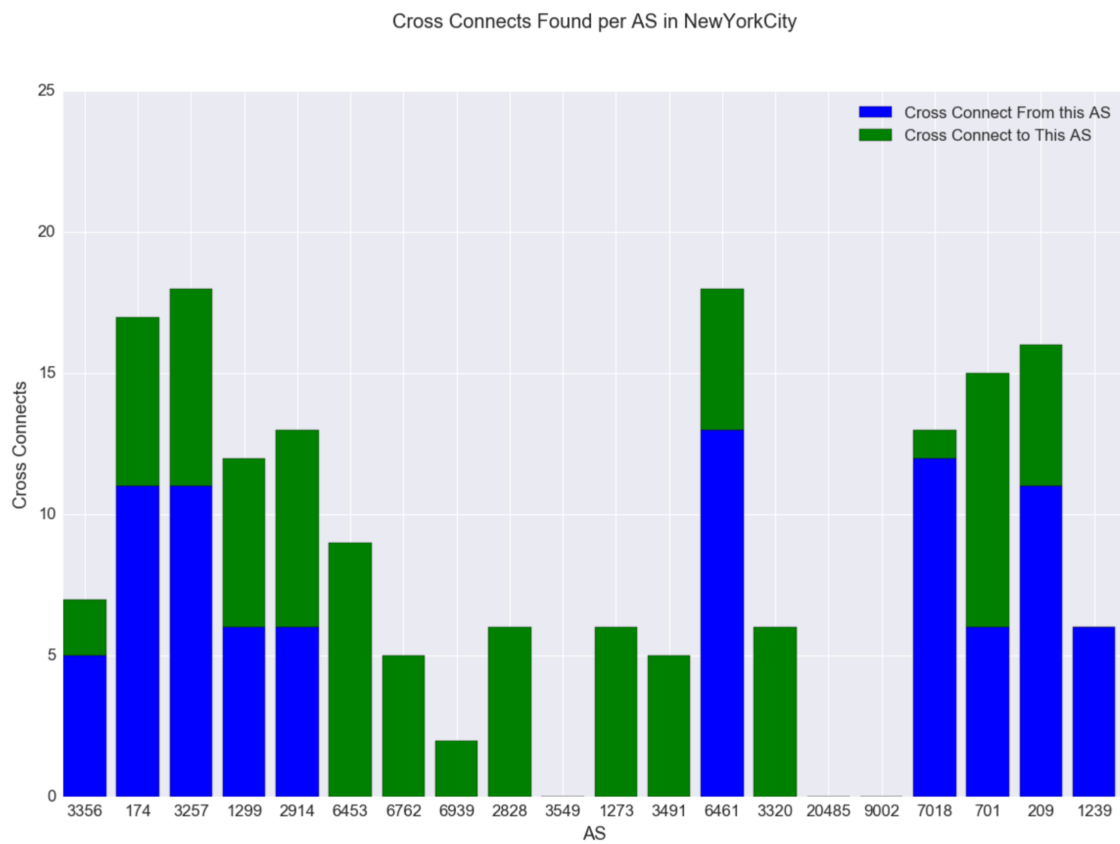

Figure 5.11: Bar chart of AS involvement in cross connects in New York City

In terms of AS presence, New York City reflects Los Angeles, with several of the ASes in figure 5.11 towering over the others. AS174, AS3257, AS6461, AS7018, AS701, and AS209 are some of the ASes that have the most involvement in cross connects. So far a trend has emerged in each of the four cities we have covered: ASes involved in the most cross connects also tend to serve as the facilitator. It seems that in each city, there are AS in the top 20 that serve as a backbone, sending data to other ASes. ASes with lower involvement in cross connects are generally recipients.

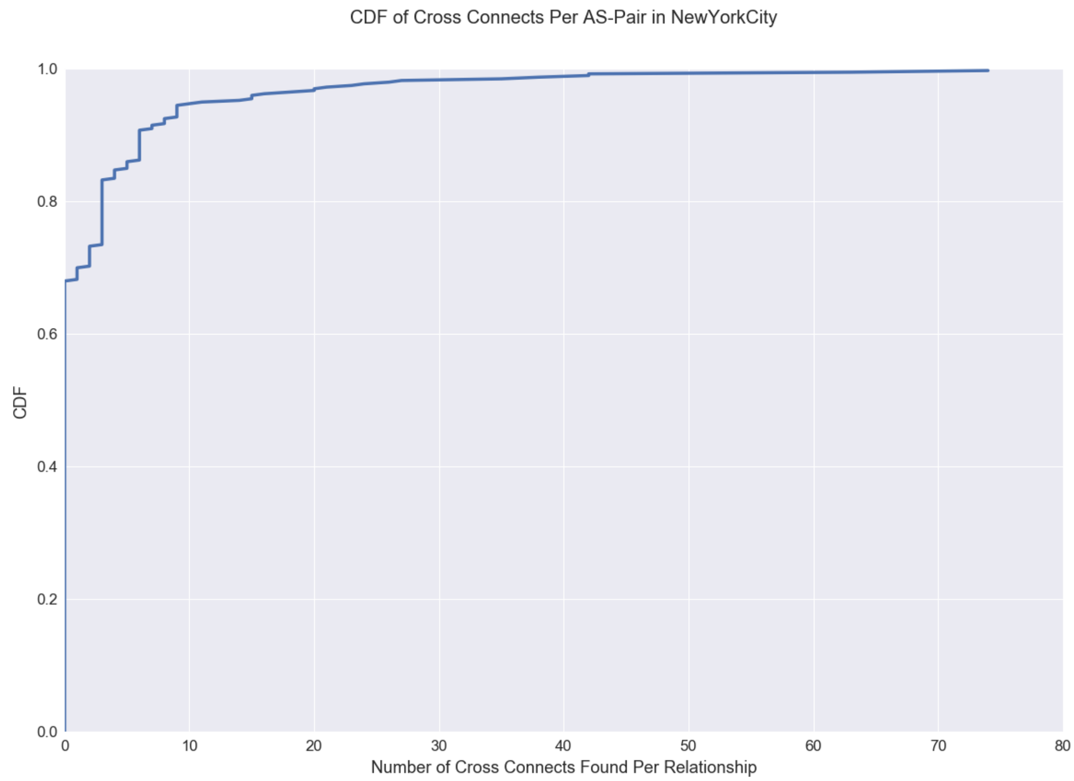

Figure 5.12: CDF of cross connects found for each AS-pair in New York City

In figure 5.12, it can be seen that New York has a group AS-pairs that have far more cross connects than the rest. Unlike the previous cities, which had about 80% of AS-pairs with no cross connect, New York City has 70% of AS-pairs with no cross connect. This reflects New York's importance, and lines up with the observations made in subsection 5.2 where it was seen that New York City utilizes 40% of potential PeeringDB connections versus 20% used by Los Angeles and Chicago. Some AS-pairs have over 70 cross connects, and it can be seen that there is an exclusive percentage of pairs that have significantly more than the rest. 25% of AS-pairs have below 10 cross connects.

Summary

- AS3356, AS174, AS3257, and AS1299 seemed to serve as ‘backbones’ for the cities examined in this section. They are involved in a large number of cross connects, and are always facilitators. Most other top 20 ASes either weren’t present, or acted as recipients.
- In Los Angeles and New York, there were several AS-pairs that had significantly more cross connects than other cities. This might be a characteristic of cities of that size.

Chapter 6: AS-Level Analysis

Based on observations made in the previous section, it is both interesting and important to investigate the behavior of ASes across the entire data set that was collected for all 17 cities. While the data collection was city focused, it is also possible to derive insights about the top 20 ASes. There are a few motivating questions for the following analysis:

1. In which cities do the top 20 ASes have the largest presence, and in which does an AS have the smallest presence?
2. Which ASes are present in the most cities?
3. Which ASes predominantly serve as facilitators with respect to the direction of the identified cross connect versus being the recipient of a cross connect?

6.1 - In Which Cities Does an AS Have the Largest/Smallest Presence?

In order to visualize this data, we will take a look at three different bar charts. The first focuses on a single AS. In this case, the AS will be AS3356, which as seen in Chapter 5, has a large presence in several cities. This chart has cities on the x-axis and the count of cross connects on the y-axis. As with some charts we have seen before, this one is split into blue and green, with blue indicating that the cross connect had the AS in question at the front, and green indicating that the AS was at the end of the cross connect.

It is apparent from figure 6.1 that AS3356 is present in a certain group of cities far more often than others. Seattle, San Francisco, Los Angeles, Dallas, and Washington D.C. all contain many more cross connects than the rest of the field. There is even a city, Columbus, that has absolutely no presence of AS3356, which is certainly surprising since the deployment of AS3356 seems to be so widespread. Why isn't AS 3356 present in Columbus?

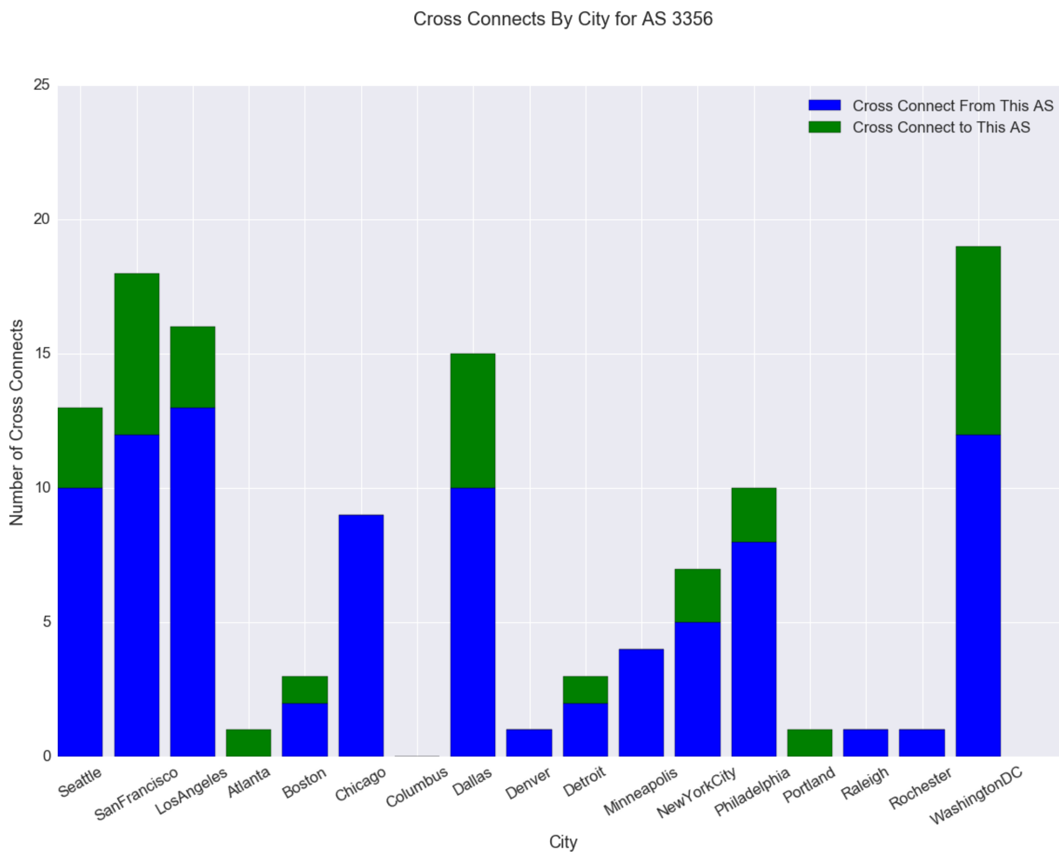

Figure 6.1: Bar chart of cross connects found in each city for AS3356

Using figure 6.2, we can take a look at the data that was collected for each city, which will give a general sense of what was found in a city. It shows the number of ASes that were found present in each city, and how that presence was determined: either through a cross connect with an RTT below the threshold (indicated by red), PeeringDB (indicated by yellow), or both (indicated by blue). This chart has cities as its x-axis, and number of ASes as its y-axis.

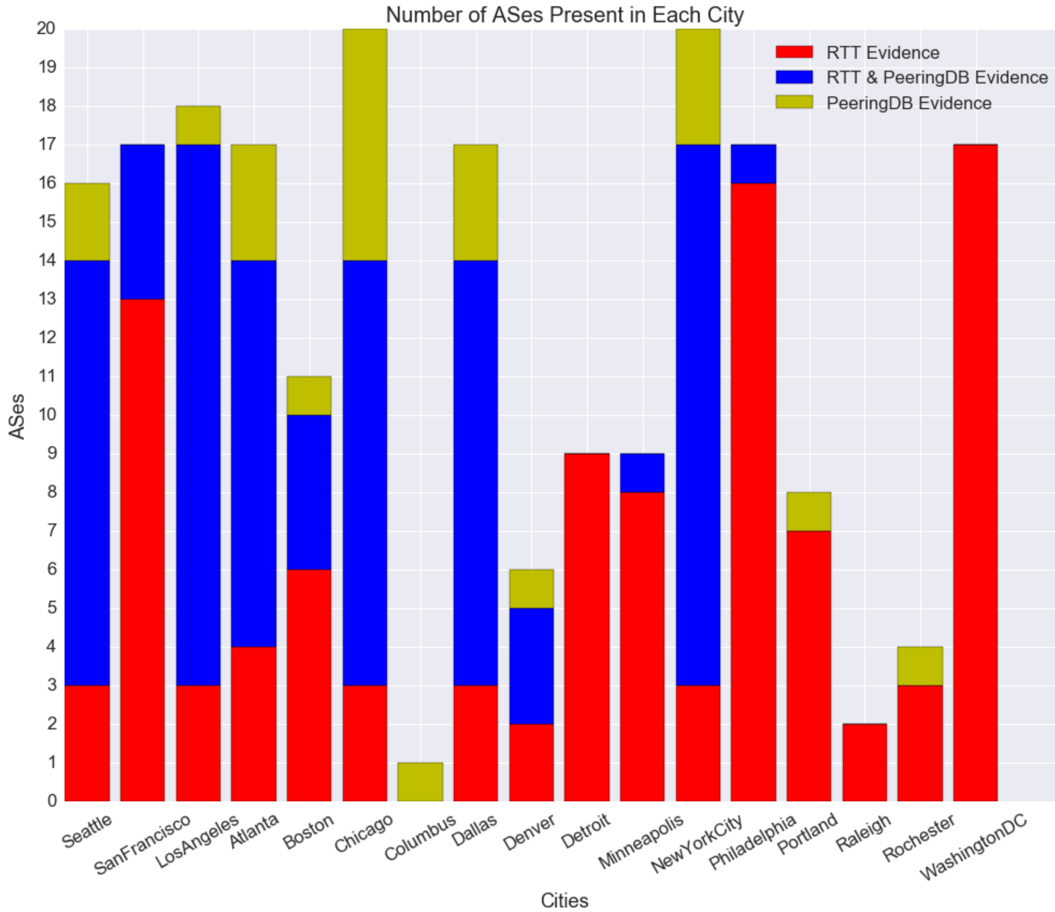

Figure 6.2: Bar chart of the number of ASes found present in each city

To return to our Columbus question, we see that Columbus was the only city in the entire study that actually had no cross connects between top 20 ASes fall below its RTT threshold. There is only evidence of a single top 20 AS being present in Columbus, and it is provided by an entry in PeeringDB. This is an indicator that Columbus was not suitable for this study. Despite the city having the requisite number of probes to be selected, it did not provide useful results.

There are several other things that we notice from this graph, however. Only two cities, Chicago and Washington D.C., had all top 20 ASes present. Furthermore, notice how the yellow bars are, for the most part, small when compared to the blue bar that they are on top of (except in the case of Columbus, of course). This means that for the majority of cases, when PeeringDB said an AS was present in a city, our methodology was able to confirm it.

While Columbus had no RTT evidence, presumably because it was a poor city to run our methodology on, Washington D.C. had absolutely no PeeringDB data regarding top 20 ASes. Why doesn't Washington D.C. have any PeeringDB data? It could be that Washington D.C. actually has no colos or IXPs in it, perhaps due to legal reasons. Another reason could be that they are registered in one of the many nearby cities. A simple online search reveals that companies Equinix and CoreSite maintain colos and IXPs in cities near Washington D.C. such as Ashburn, VA and Reston VA [15, 16].

Summary

- It is possible to look at the presence of an AS within the cities surveyed. In this case, the city presence of AS3356 was visualized. Certain cities had disproportionately more cross connects within them (Dallas, Seattle, San Francisco, Los Angeles, and Washington D.C.). AS3356 had no cross connects in Columbus
- A bar chart of the number of ASes found in each city shows that Columbus had no RTT confirmed AS presence whatsoever, and only one AS present according to PeeringDB. This calls into question Columbus' suitability for this study.

- When PeeringDB indicates the presence of an AS, our methodology was for the most part able to confirm it.
- Washington D.C. has no PeeringDB data regarding AS presence.

6.2 - Which ASes Are Present in the Most Cities?

There is another angle that should be investigated. We took a look at how many ASes are present in a city, but it is also important to take a look at how many cities a particular AS is present in. Figure 6.3 has ASes as its x-axis, and number of cities for the y-axis. It is color coded in the same way as the previous chart.

There are a few notable observations to be made here. First, AS3356 is present in the most cities, at 16. This AS, alongside AS174 and AS7018, are also the only ASes without any PeeringDB data. Therefore, we use only RTT based evidence regarding their presence within a city. It’s extremely unlikely that the two largest ASes (and another very large AS in AS7018) aren’t utilizing colos or IXPs, so this lack of PeeringDB data suggests that these three ASes aren’t registering their presence with PeeringDB.

Next, much like figure 6.2 in the previous section, it can be seen that blue makes up a large part of most bars. This once again shows that our methodology and PeeringDB corroborate each other rather well. However, there are two ASes, AS20485 and AS9002, that have no RTT data and have only PeeringDB data indicating their presence in cities. Furthermore, as can be seen in Chapter 5, AS20485 and AS9002 were never found to be involved in any cross connects below the RTT threshold for any of the cities investigated.

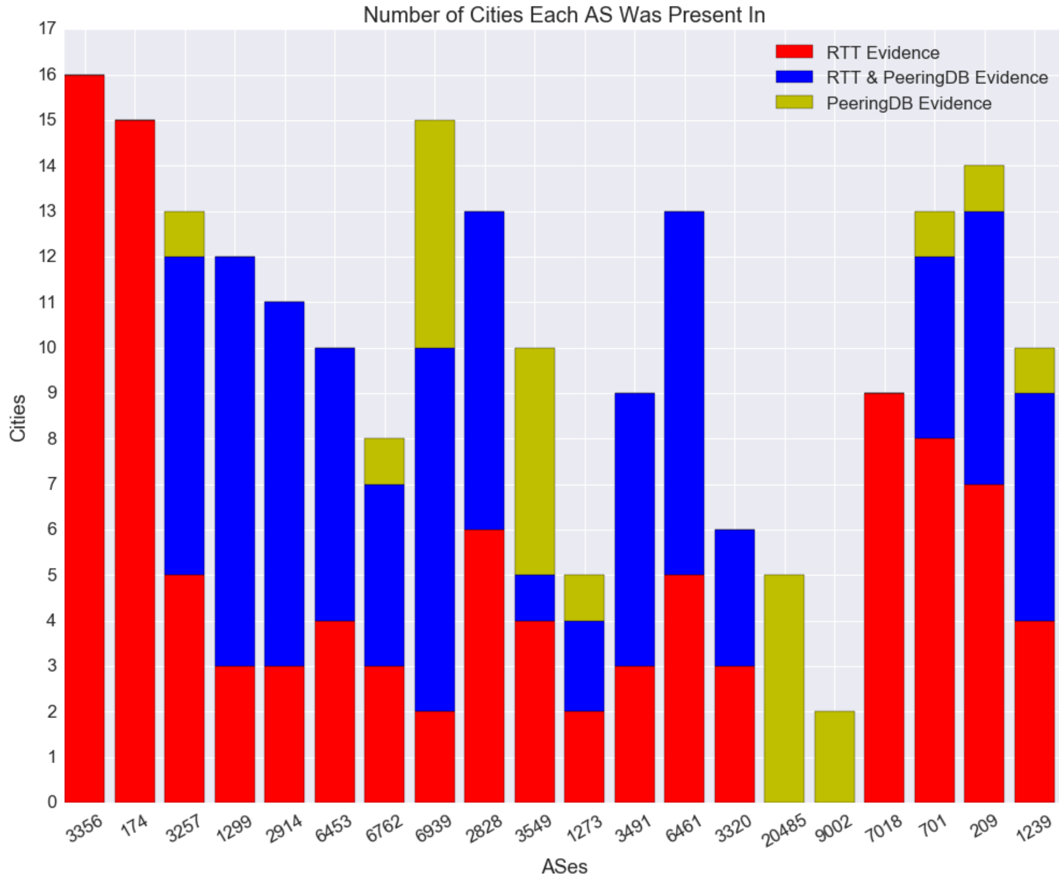

Figure 6.3: Bar chart of the number of cities each AS was present in

Not only that, but if one looks at the AS matrices in Chapter 5, it can be seen that AS20485 and AS9002 were never found to be involved in a single cross connect, even one above the RTT threshold. Why is this? Upon further investigation, it can be seen that despite the presence of AS20485 and AS9002 in the colos and IXPs of several cities in the United States, these are foreign ASes that are based in Eastern Europe and Russia. Therefore, while they have a presence in facilities within the United States, unless we sent a traceroute to a destination in Russia, the chances of a cross connect being found with these ASes was slim. As such, despite the sheer number of traceroutes launched, not a single cross connect with AS20485 or AS9002 was found.

Summary

- No ASes were found to be present in all 17 cities.
- AS3356, AS174, and AS7018 do not have PeeringDB data to indicate AS presence.
- AS20485 and AS9002 have a presence in several United States colocation facilities and IXPs, but were never involved in any cross connects. This is most likely because they are based in Eastern Europe and Russia.

6.3 - Which ASes Predominantly Serve as Facilitators, and Which Are Recipients?

In this section, we will use a data visualization that allows us to see which top 20 ASes a particular AS connects to, while also indicating the nature of the relationship. The x-axis has all of the top 20 ASes, and the y-axis shows the number of cross connects found. Once again, green indicates that the connection was going towards the AS, and blue indicates it was leaving that AS.

In order to see how this chart can show us whether or not an AS is a facilitator versus a recipient of network traffic, we will take a look at two examples. The first is the chart for AS3356, shown in figure 6.4. Looking at the chart, one can observe that for each relationship, AS3356 participates in a majority of its cross connects as the facilitator, sending the data to the other AS. For some ASes, the relationship is completely one-sided, with all cross connects going from AS3356 to the other AS.

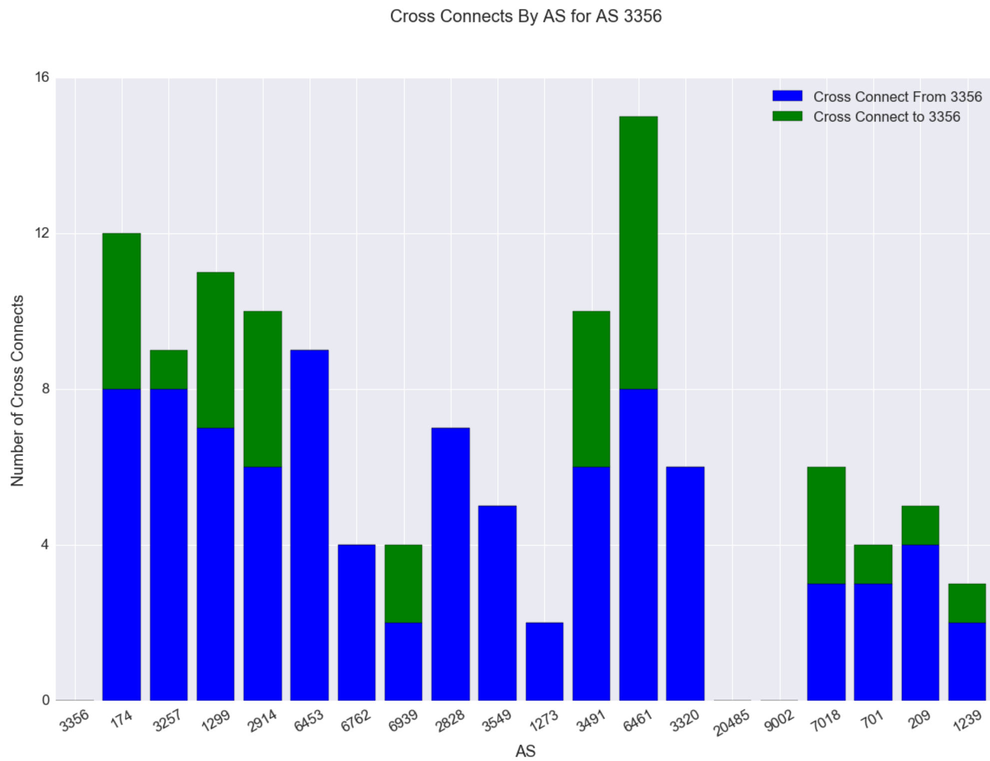

Figure 6.4: Bar chart of the number of cross connects AS3356 has with top 20 ASes

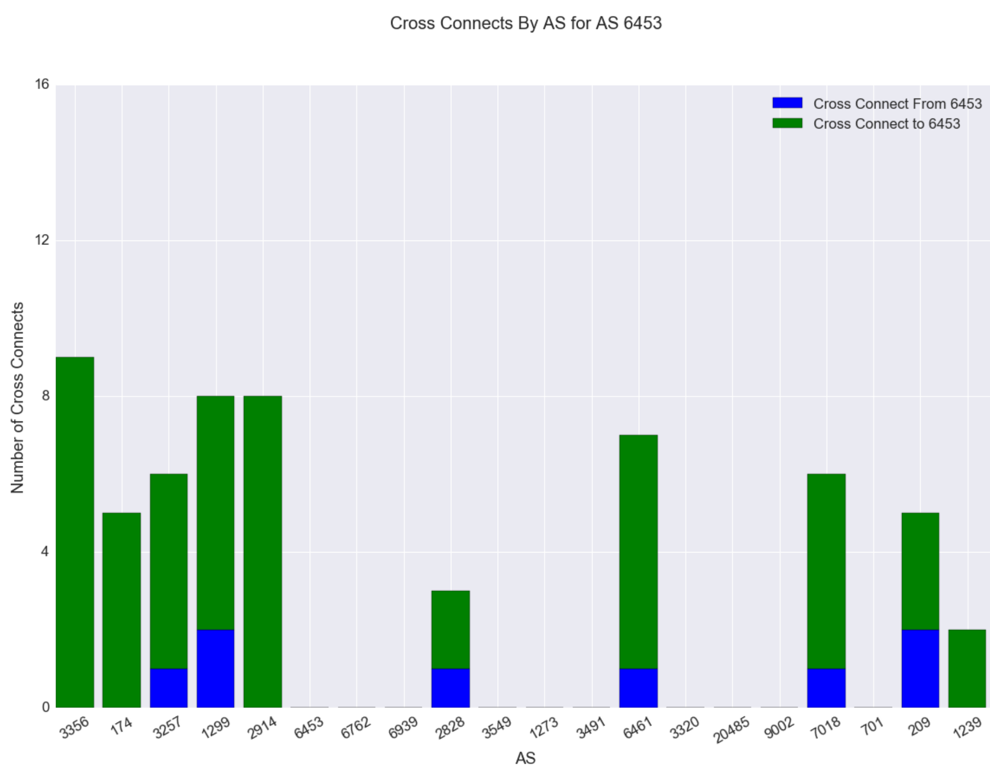

Figure 6.5: Bar chart of the number of cross connects AS6453 has with top 20 ASes

The next chart is of AS6453, shown in figure 6.5. This is shown to contrast with AS3356. This AS can be characterized as a recipient according to our data, since it has relationships that generally flow in one direction. With AS3356 in particular, all cross connects only send data towards AS6453, showing that the relationship is predominantly one-directional. It also doesn't have as many cross connects with as many different ASes as AS3356, which reflects its smaller size.

Summary

- Using the data, it is possible characterize ASes as either a facilitator or a recipient by observing the direction of their relationships with other ASes.
- AS3356 is a facilitator since it is most frequently found at the front of a cross connect with other ASes.
- AS6453 is a recipient since it is most frequently found at the end of a cross connect with other ASes.

Chapter 7: AS-Pair-Level Analysis

The final layer of analysis investigates the AS-pair. Since this study focuses on the top 20 ASes, there are 190 AS-pairs. These pairs won't necessarily have a cross connect found between them. In fact, as will be seen, many AS-pairs did not have a cross connect.

Here are the motivating questions:

1. Do ASes favor some cities more in terms of the number of cross connects that they have present there?
2. Is there a pattern regarding where cross connects between AS-pairs are located?
3. Of the ASes present in a city, what percentage of possible AS-pairs have a cross connect between them?

7.1 - Do ASes Favor Some Cities More in Terms of the Number of Cross Connects That They Have Present in a City?

The following chart shows the number of cities that an AS was present in through the transparent red bars. On top is a scatter of points, with each representing a city that the AS is present in. The height of the point shows the number of cross connects that were present in the city it represents. The x-axis contains the names of the top 20 ASes, while the y-axis measures the count for both the bar chart and the scatter plot. This chart conveys a sense of how distributed an AS is. If the scatter is compacted, then the AS has a similar presence in multiple cities. Any outliers indicate a city in which the AS had a disproportionately large or small footprint compared to the rest of the cities it was in.

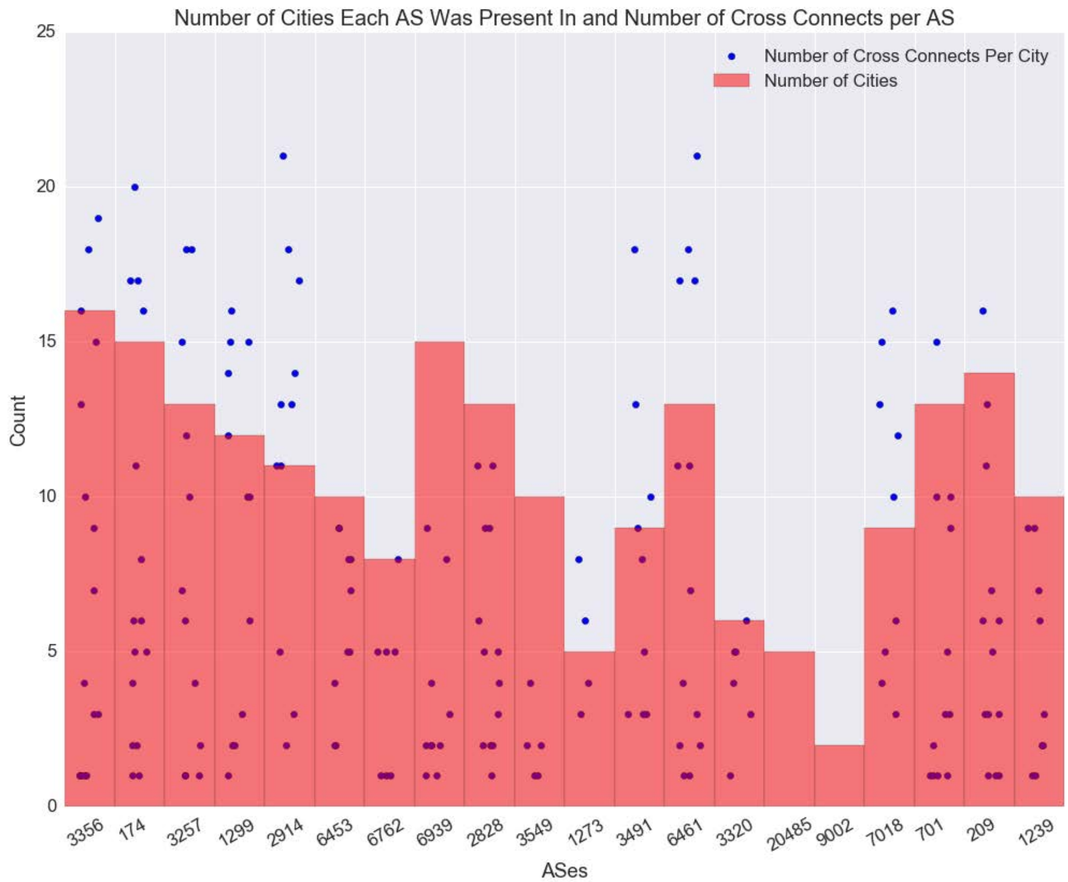

Figure 7.1: Bar chart showing number of cities each AS was present in on top of a scatter plot of cross connects per city per AS

What is clear is that in almost all cases, the number of cross connects in cities varied. For example, AS2914 can be seen having connections in 11 cities. 3 of these cities had below 5 cross connects, and one had over 20.

Summary

- The number of cross connects per city for a given AS varies, with some cities exhibiting many more cross connects for a given AS.

7.2 - Is There a Pattern Regarding Where Cross Connects Between AS-Pairs Are Located?

The observations of the previous section naturally lead to this question. It is clear that ASes have more cross connects in some cities than others. Is this true for relationships as well? Groups of cities with the largest numbers of common AS-pairs with cross connects were aggregated. Table 7.1 summarizes the results. The table shows each group of size 2 through 9 with the largest number of common cross connects.

Number of Cities	Cities	Number of Common AS-Pair Cross Connects
2	New York City, Philadelphia	55
3	New York City, San Francisco, Philadelphia	40
4	New York City, San Francisco, Washington D.C., Philadelphia	28
5	New York City, San Francisco, Dallas, Philadelphia, Los Angeles	18
6	New York City, San Francisco, Dallas, Philadelphia, Washington D.C., Los Angeles	15
7	San Francisco, Dallas, Washington D.C., New York City, Philadelphia, Los Angeles, Seattle	12
8	San Francisco, Dallas, Washington D.C., New York City, Seattle, Chicago, Philadelphia, Los Angeles	9
9	San Francisco, Dallas, Seattle, New York City, Chicago, Atlanta, Washington D.C., Philadelphia, Los Angeles	4

Table 7.1: City groups with the most common cross connect counts for groups of cities size 2 to 9

This table show that there are clearly some cities that top 20 ASes use the most for connecting with other top 20 ASes. In the group of nine, the cities involved create a coverage of the United States, with cities in every major region of the United States. Some cities, such as New York, Philadelphia, and San Francisco, appear in almost every group, indicating that they are important cities for common AS-pair connections. Each of these cities are large cities with well-developed infrastructures, which may contribute to why ASes choose to connect in those cities.

Summary

- San Francisco, New York City, and Philadelphia are popular cities for ASes to cross connect.
- Larger cities tend to be more involved when searching for common cross connects amongst a group of cities.

7.3 - Of the ASes Present in a City, What Percentage of Possible AS-Pairs Have a Cross Connect Between Them?

The following charts are CDFs that visualize results regarding the 190 possible AS-pairs in this study. The first shows the number of cross connects found per AS-pair. There are several observations that can be taken from figure 7.2. First, over 40% of the 190 AS-pairs possible had no cross connect found between them. The next approximately 30% of AS-pairs had between 1 and 5, and the final 30% had between 6 and 15 cross connects. Thus, most AS-pairs that have a cross connect had just several cross connects found in all of the cities surveyed. There are a few AS-pairs that show up much more often, most likely involving some of the ASes that we have seen have the largest national presence, such as AS3356. Figure 7.3 shows the number of cities that AS-pairs were found in. Since 40% of AS-pairs didn’t have a single cross connect, 40% didn’t show up in any cities.

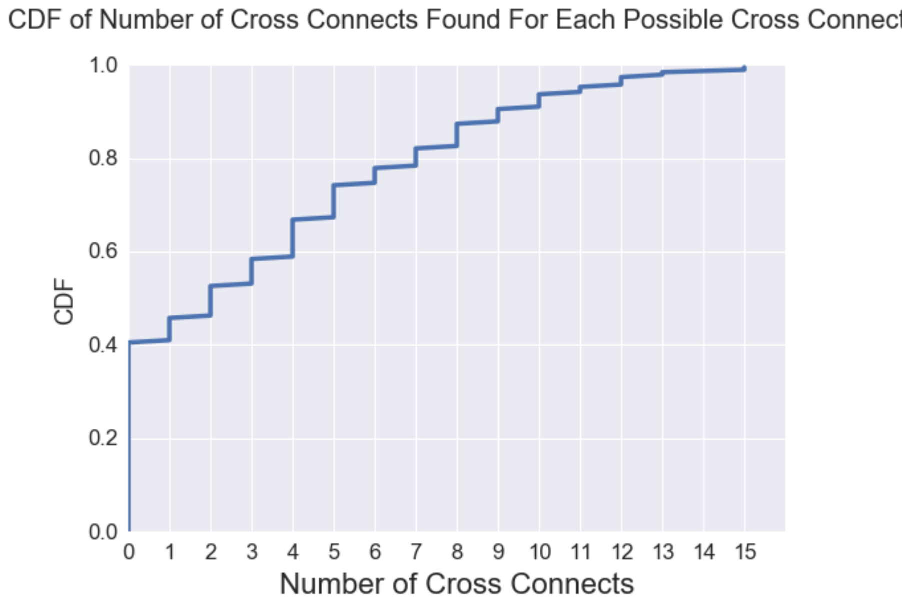

Figure 7.2: CDF showing the number of cross connects found for each AS-pair

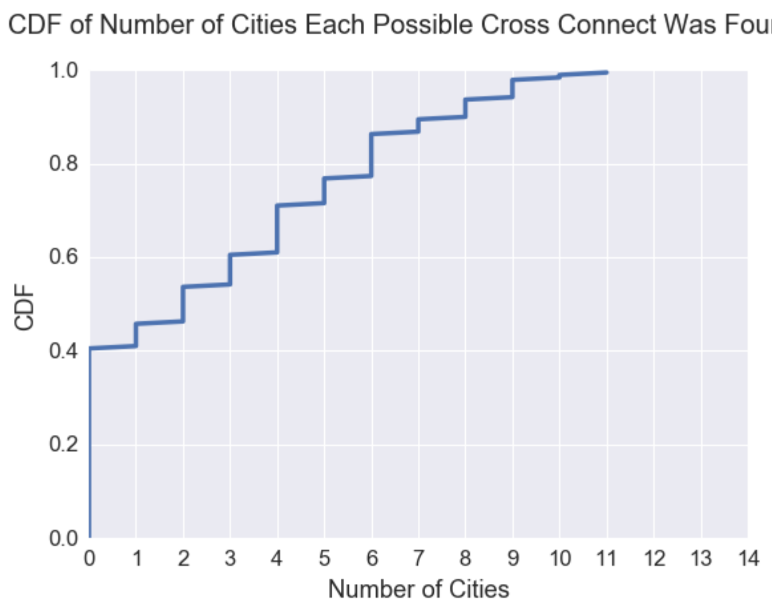

Figure 7.3: CDF showing the number of cities each AS-pair was present in

There aren't really many features of note in figure 7.3, as the distribution is fairly linear. After 40%, the first 30% of AS-pairs appeared in 1 to 5 cities. After that, the final 30% show up in 6 to 11. There isn't any pattern or bulge or irregularity that suggests some sort of a quirk. There are very few AS-pairs present in 10 or 11 cities, but other than that, the distribution is fairly even.

Summary

- There are a certain top 5% or so of AS-pairs that had more AS-pairs found throughout the country. There are most likely pairs that involve ASes with a large national presence, such as AS3356.
- Though a few AS-pairs were found in 10 to 11 cities, there didn't seem to be any AS-pairs that had a outlying number of cities in which they were present.

Chapter 8: Related Work

This study does not seek to redefine the state of the art for determining the geolocation of Internet infrastructure. Rather, it is a small piece in understanding the benefits and limitations of the methodology used in this thesis, and is in fact a part of a larger NSF funded project called “NeTS: Small: Towards an Accurate, Geo-Aware, PoP-Level Perspective of the Internet’s Inter-AS Connectivity.” [18] This project seeks to develop a methodology for determining the location of AS cross connects and AS points of presence at a geographic level by both determining where ASes connect independently and where ASes connect within colocation facilities or IXPs. The main goal of the project is to design, develop and evaluate techniques to map the geographic location of a given target AS and determine the inter-AS connections of an AS. In short, we are interested in producing for any given AS a corresponding map of its presence and cross connects at the city-level.

The existing body of related work falls into a few different categories, as the topic of IP geolocation can itself be broken into several different areas: colocation facilities/IXPs, core IP addresses, and edge IP addresses. The geolocation of each provides its own challenges. This thesis focuses on the first two areas, since cross connects by their definition cannot occur at the edge. Furthermore, ASes of this size generally do not exist at the network edge.

There are numerous techniques that have been investigated and developed (some by ONRG researchers) in order to determine the geographic location of cross connects [6]. Some are limited to colos and IXPs, but others also attempt to find the geolocation of independent routers. This thesis investigated the use of RTT thresholds to determine where these cross connects occur at different cities. Other studies have also used traceroute measurements to determine algorithmically where cross connects are physically located. One utilized Border Gateway Protocol data and DNS clues to infer existing AS-level links at colocation facilities, but did not utilize RTT as a method of determining location, nor did it utilize RIPE Atlas as a platform [1].

There has also been work that showed the relationship between distance that a traceroute has to travel geographically and the distance error found in measurements calculated by algorithms that utilize delay. It showed that delay works well as an indicator of locations when the traceroute is launched near the host being targeted [8].

Another study focused on edge hosts in order to characterize the size of customer facing, or “eyeball”, ASes. According to this paper, determining the geolocation of hosts in the network core is considered more difficult than inferring the location of an edge host [6].

The results of this study need to be combined with the methods of other studies that have already been conducted to help corroborate these results and determine if RIPE Atlas in conjunction with RTT distribution is an effective geolocation tool. As far as we are aware, using the distribution of RTT to determine threshold RTTs, much less threshold RTTs in general, have not been used to estimate the location of a router at a city-level. While the results present in this study seem to suggest a correlation between these RTT thresholds and AS location according to PeeringDB, it is clear that more support is necessary. Utilizing the methods in these other studies could prove critical to eventually utilizing RTT thresholds in large scale traceroute campaigns for effective city-level AS cross connect geolocation.

Chapter 9: Conclusions and Future Work

This work is part of a larger whole, and only begins to answer the questions of cross connect geolocation. In some cities, such as Chicago, using a distribution of RTT values yielded convincing results, with a clear separation between two distinct zones. Presumably, these zones are the area within the city and the area outside of it. The data for some cities is not as clear, and there must be further study to understand what specifically made Chicago's campaign so successful. If the conditions and parameters are replicable in other cities, then the analysis of those cities can be improved upon for future study. However, that is not to say that the campaigns done for other cities were not useful.

There are conclusions to be gleaned from the campaigns launched even in cities like Dallas, which had very poor distributions of RTT values that were not conducive to RTT threshold estimations. For those cities, the discussion revolves around the observation that a majority of the ASes that PeeringDB states are in the same colo or IXP don't seem to take advantage of the opportunity to interconnect. The low rate of traceroutes passing through potential AS-pairs as reported by PeeringDB implies a low rate of top 20 ASes establishing cross connects in colos and IXPs. In terms of further research, there must be an investigation done to see exactly what large ASes are doing in colos and IXPs in terms of connectivity. Large ASes aren't going to be paying fees to the owner a facility to simply sit idle. It might be possible that large ASes use colos and IXPs to connect with smaller ASes. The methods included in a couple other studies [1, 6] would be invaluable to pursuing this question.

Those methods would also be useful to help confirm the actual presence of these cross connects within the surveyed city. What this study lacks are sufficient confirmation methods for the RTT threshold. PeeringDB alone is not enough to corroborate the presence of an AS cross connect within a city, since direct links are also possible. We believe that a combination of geolocation methods utilized in several other papers could help establish this as a reliable method for determining AS cross connects [1, 8, 6, 17, 11].

An observation to be taken from the AS-level view are the directional relationships present between ASes. Certain ASes such as AS3356 act as facilitators of data, and sometimes have one-way relationships with other top 20 ASes. The business relationships between ASes can be inferred by viewing a graph of the outgoing and incoming communication of each AS in conjunction with some context about the company that owns the AS.

In summary, there are a few things to take away: under specific (currently unknown) conditions, a distribution of RTT values can provide a clear distinction between two different RTT ranges. Since RTT is in part related to distance [17], it is a reasonable conclusion that one of these regions pertains to some region of a city, whereas the other pertains to some region outside of it. Next, there seems to be a shortage of utilized opportunities for the top 20 ASes to connect within colos and IXPs within which they are present. Finally, the direction of AS cross connects can be utilized to infer high level views of AS customer and business relationships.

Topics for further study must inquire into what the conditions are that made a city such as Chicago have such clear cut division between their RTT ranges. This will be critical moving forward to develop this method, in conjunction with added confirmation methods to provide a clearer conclusion as to whether or not these cross connects are in the city. Further study must also be devoted to understanding why top 20 ASes seem to be underutilizing connections that are available in colocation facilities and IXPs, if that is indeed the case. For that, there needs to be additional investigation into how to confirm a connection between two ASes in a facility, which has already been done [1] but needs to be applied in the context of this study. An immediate application of this data is a visualization tool that allows a user to view AS cross connect and AS presence data on a map.

Chapter 10: References

[1] Motamedi, R., Chandrasekaran, B., Maggs, B., Rejaie, R., & Willinger, W. (n.d.). **On the Geography of X-Connects.** *In Progress*

[2] **AS Rank: AS Ranking - CAIDA**: http://as-rank.caida.org. (n.d.). Retrieved May 1, 2016, from http://as-rank.caida.org

[3] Hawkinson, J, and Bates, T. **Guidelines for creation, selection, and registration of an Autonomous System (AS).** RFC 1930 (Best Current Practice), Internet Engineering Task Force, Mar. 1996, updated by RFC 6996

[4] Kurose, J., & Ross, K. **Computer networking: A top-down approach (4th ed.).** Boston: Pearson/Addison Wesley, 2008

[5] L. Gao, **"On inferring autonomous system relationships in the Internet,"** IEEE/ACM Transactions on Networking (ToN), vol. 9, no. 6, pp. 733–745, 2001

[6] Motamedi, R., & Willinger, W. (n.d.). **A Survey of Techniques for Internet Topology Discovery.** *IEEE Communications Surveys & Tutorials IEEE Commun. Surv. Tutorials*, 1044-1065

[7] **Team Cymru: IP to ASN Mapping.** http://www.team-cymru.org/IP-ASN-mapping.html#whois, 2015. Retrieved May 17, 2016

[8] Katz-Bassett, E., John, J., Krishnamurthy, A., Wetherall, D., Anderson, T., & Chawathe, Y. (n.d.). **Towards IP geolocation using delay and topology measurements.** *Proceedings of the 6th ACM SIGCOMM on Internet Measurement - IMC '06*

[9] **IP2Location.** http://www.ip2location.com, 2016. Retrieved May 17, 2016

[10] **MaxMind.** https://www.maxmind.com/en/home, 2016. Retrieved May 17, 2016

[11] Rasti, A., Magharei, N., Rejaie, R., & Willinger, W. (n.d.). **Eyeball ASes: From Geography to Connectivity**. Proceedings of the 10th Annual Conference on Internet Measurement - IMC '10

[12] **Akamai EdgeScape Service Description.** http://goo.gl/akSvGu, March 2002. Retrieved May 1, 2016

[13] **PeeringDB.** https://www.peeringdb.com/, 2016. Retrieved May 17, 2016

[14] **RIPE.** http://ripe.net, 2011. Retrieved November 11, 2015

[15] **Equinix DC1 Washington DC – Network Interconnection Specifications.** http://www.equinix.com/locations/united-states-colocation/washington-dc-data-centers/dc1, 2016. Retrieved May 16, 2016

[16] **CoreSite Washington D.C. & Northern Virginia Data Centers.** http://www.coresite.com/data-centers/locations/northern-virginia-washington-dc/dc1, 2016. Retrieved May 16, 2016

[17] Padmanabhan, V., & Subramanian, L. (2001). **An investigation of geographic mapping techniques for internet hosts.** *ACM SIGCOMM Computer Communication Review SIGCOMM Comput. Commun. Rev.*, 173-185

[18] Rejaie, R., Willinger, W., & Maggs, B. **NeTS: Small: Towards and Accurate, Geo-Aware, PoP-Level Perpsective of the Internet's Inter-AS Connectivity.** NSF Proposal

[19] Durairajan, R., Barford, P., Sommers, J., & Willinger, W. **InterTubes: A Study of the US Long-haul Fiber-optic Infrastructure.** SIGCOMM '15, August 17–21, 2015, London, United Kingdom