27 November 2014

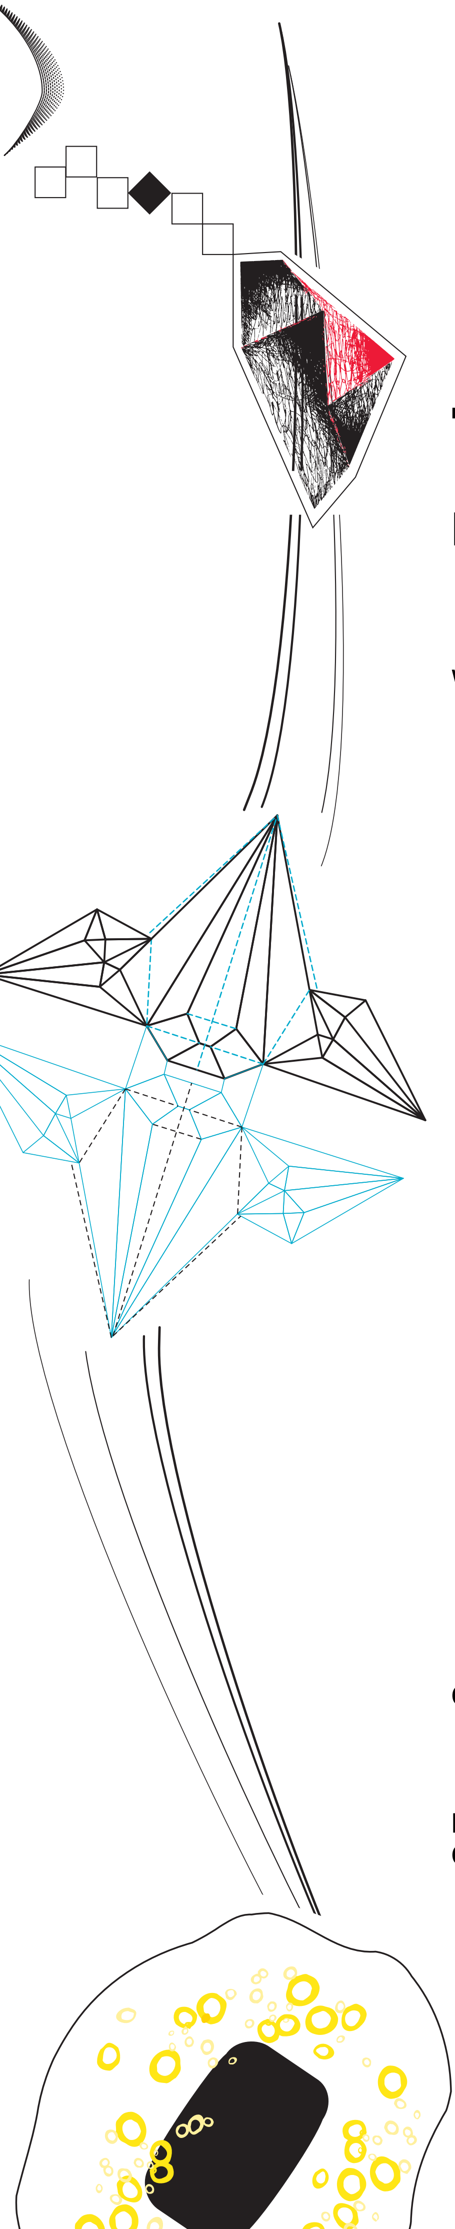

HOW ASYMMETRIC IS THE INTERNET?

MASTER THESIS

Wouter de Vries, s1023403

Committee: prof. dr. ir. Aiko Pras
dr. Anna Sperotto
José Jair Santanna, M.Sc

Faculty: Electrical Engineering, Mathematics and Computer Science (EWI)
Chair: Design and Analysis of Communication Systems (DACS)

UNIVERSITY OF TWENTE.

CONTENTS

CHAPTER 1

INTRODUCTION

This document is the final product of 9 months of work at the Design and Analysis of Communications Systems (DACS) group at the University of Twente. As opposed to a traditional thesis, this one is built around a paper that will be submitted to one of the main European network measurement conferences, the 7th Traffic Measurements and Analysis Workshop, or TMA 2015[1]. TMA 2015 is a conference on almost all types of network measurements. Aside from the paper the research was also presented at a meeting that was organized by Reseaux IP Europeens (RIPE), called RIPE69[2]. The main goal of this document is to present the paper that was written during the Master Assignment and how that work meets the requirements as defined in [23]. In addition, this document serves to show the additional work that has been done.

The structure of this document is as follows: First, the Master requirements will be introduced and how they are fulfilled in Section 1.1. Then, in Section 1.2, a short summary of the research topics that were written prior to the final project. The deliverables, such as the paper, will be described in Chapter 2. The paper itself and the slides that were used for the RIPE69 meeting are included in Appendix A and B.

1.1 Master requirements

The Master Assignment is not usually finished by writing a paper. Therefore, this section is aimed at making it clear how the Master requirements, as described in [23], are met.

1.1.1 Formulate a problem statement

As part of the research topics four research questions, described in subsection 1.2.2, were formulated regarding Distributed Denial of Service (DDoS) Mitigation. The paper that was written for TMA 2015 had one of these research questions, RQ3, as its basis. The first two research questions, RQ1 and RQ2, were answered during the initial literature study. However, during the study the third research question (RQ3) turned out to require a large scale investigation. As answering RQ4 was dependent on RQ3 we chose to focus solely on that one. Formulating a problem statement is an essential part of writing a paper, for that reason, I spent a lot of time on clearly defining the problem in the paper and thus meet this requirement.

1.1.2 Identify relevant literature

The start of every Master Assignment within the Computer Science department is Research Topics. This consists of identifying relevant literature and describing your findings in a document, which is summarized in Section 1.2. While that already satisfies the requirement, relevant literature is also identified and cited, where necessary, throughout the paper and in the presentation.

1.1.3 Draw up a work plan

This requirement was fulfilled in two ways. First, during the Research Topics the approach for answering the research questions was determined along with a preliminary schedule. The word ’preliminary’ implies that the final schedule was different, as is indeed the case. The main change to the schedule was instigated by the change of focus from DDoS mitigation to Internet routing asymmetry. Second, the paper was planned according to the deadline of the conference that it will be submitted to.

1.1.4 Adjust goals and approaches based on interim evaluations

The goals and approaches of this study were altered several times. The best example of this is the fact that the goal of this assignment changed from implementing a DDoS mitigation solution to analyzing the

[1] http://tma-2015.cba.upc.edu/

[2] https://ripe69.ripe.net/

characteristics of Internet routing asymmetry. This change was instigated almost halfway through the assignment based on interim evaluations with my daily supervisor.

1.1.5 Analyse different possible solutions and motivate a choice between them

Various methods for doing the measurements that were needed to answer the research questions were considered, the choice for one of these was motivated in the paper. Also, several methods of analyzing the results were considered and chosen based on meetings with other group members and/or my daily supervisor.

1.1.6 The ability to reflect on the problem, on the research/design approach, on the solution and on ones own performance

The problem that was first considered, the mitigation of Distributed Reflection Denial-of-Service attacks, is not the one that ended up being answered. The reason for this is that while evaluating the approach it became clear that a solution to another problem had to be found first. After deliberation with my daily supervisor I decided to focus on that new problem. Furthermore, part of writing the paper was finding the weak points of the study so that they could either be repaired or at least be noted as a point of caution to readers.

Regarding my own performance, most people who are working on their final Master Assignment get to a point where they feel as if their research is going nowhere, I was no exception in this. During the assignment I had to look at my own work and my performance frequently to determine whether certain goals were still feasible, and sometimes this turned out not to be the case so they had to be adjusted.

1.1.7 Demonstrate creativity and the ability to work independently

After collecting the data from the measurements I first had to come up with some way to analyze the data. The large amount and the nature of it required the design of a new analysis tool. This required creativity for two reasons. First, no one had done asymmetry measurements on this scale before, or using this method. Second, no one had done the analysis that I wanted to do before. Aside from the actual analysis of the data I had to come up with various ways to present it to others. The analysis itself was done largely independently from my supervisor, however, we regularly conferred on the direction of the study.

1.1.8 Communicate the research and design activities both written and in presentations

Writing the paper for TMA 2015 required writing down what had been done in a clear and concise way. The sole goal of the paper is to communicate the results of the study with the rest of the world which is in clear agreement with this requirement. In addition, the research and the results were accepted to be presented by me at a meeting organized by the Reseaux IP Europeens (RIPE) organization named RIPE69. This bi-annual meeting, where Internet Service Providers (ISPs), network operators and other interested parties are invited to attend, covers almost all network-related topics. The research was presented in front of more than 500 people. Some preliminary results were presented to the group (DACS) to introduce them to my research and the final result of this Master Assignment will be presented at my final presentation on the 27th of November 2014.

1.2 Research Topics

In this section I will briefly summarize the literature study that was conducted prior to the start of the Master Assignment. This section is structured as follows: First, the topic will be introduced. Then, the research questions that were determined at the start of the study. Finally, the three main subjects of the literature study will be summarized, namely, the Internet, types of attack and types of mitigation.

1.2.1 Introduction

Distributed Denial-of-Service (DDoS) attacks are a constant threat on the Internet. They cause a lot of damage by interrupting services and are very difficult to defend against. DDoS attacks are a large risk for all users on the Internet as everyone is vulnerable, including large corporations [8, 13, 20, 1]. The reason for this is that DDoS attacks can exhaust the victim's available bandwidth without relying on any specific vulnerability. The amount and intensity of the attacks are on the rise [18, 7] while at the same time the number of users that rely on online services is increasing. One of the most powerful types of DDoS, Distributed Reflection Denial-of-Service (DRDoS) relies on session-less services such as the Domain Name Service (DNS) and Network Time Protocol (NTP) to amplify attacks. The amplification that these services provide result in a dramatic increase in the strength of DDoS attacks. Examples of such attacks are given in [17] and [16].

1.2.2 Research questions

As part of the research topics the following four research questions were defined:

- **RQ1:** What are the characteristics of D(R)DoS attacks?
- **RQ2:** What are the existent solutions against D(R)DoS attacks?
- **RQ3:** How to quickly identify the networks that the DRDoS attack passes through?
- **RQ4:** How to combine the networks that the attack passes through in a mitigation solution?

RQ1 and RQ2 are answered in subsection 1.2.4 and 1.2.5. RQ3 is partially answered in subsection 1.2.3.

1.2.3 The Internet

The Internet consists of a large amount of inter-connected networks. Each network consists of one or more routers which route incoming packets that are received towards their destination and/or to another router. The networks, that are identified by Autonomous System (AS) Numbers (ASN), that make up the Internet can be divided into three categories:

1. Tier 1 network, a network which peers (i.e. connects without paying for bandwidth) exclusively with other Tier 1 networks.
2. Tier 2, a network which has a transit agreement (i.e. connects and pays for bandwidth) with Tier 1 networks. In addition this type of network can have peering with other networks which are not Tier 1.
3. Other (Unofficially Tier 3), a network which exclusively uses transit with other networks and has no peering agreements.

There are only a few Tier 1 providers worldwide, while there are many more Tier 2 and Other (Tier 3) providers [24]. Most if not all routers and end-hosts on the Internet have an Internet Protocol (IP) address which uniquely identifies them on the Internet. This IP address can be used to send packets to or from and is the basis for almost all communication over the Internet. One key aspect of Internet Routing is that it is asymmetric [6, 5, 22]. It is important to note that this makes it difficult to determine how a packet reaches a destination from the viewpoint of the receiver.

1.2.4 Types of attack

There are several levels on which a (D)DoS attack can occur [9]. Examples are the application level [2, 19], operating system level [11] or protocol level [21]. The most difficult to mitigate type of DDoS attack are volumetric attacks because they do not rely on a fault in the target's systems. Volumetric means that the attack relies on exhausting the Internet bandwidth that the target has available [3]. The means of attack vary from simple flooding [9] to advanced reflection attacks, where intermediate services are exploited to amplify the scale of the attack.

1.2.5 Types of mitigation

According to [15] we can divide DDoS mitigation into four categories: prevention, detection, source identification and reaction. Of these four categories we will briefly summarize prevention, detection and reaction. Prevention means stopping attacks completely before they can be started. This method of mitigation is mostly academic and consists of methods of preventing source address spoofing. Examples of this are Best Current Practice (BCP) 38 [4], Route-based Distributed Packet filtering (DPF) [14] and Source Address Validity Enforcement (SAVE) [10]. Determining whether you are under attack is called detection. There are two basic ways this can be done [12], namely pattern detection and anomaly detection. In terms of response methods, three common approaches exist [12]. These are rate-limiting, filtering and reconfiguration. The latter being the one being used by services such as CloudFlare[3] and Prolexic[4].

[3] https://www.cloudflare.com/

[4] http://www.prolexic.com/

CHAPTER 2

DELIVERABLES

Aside from a paper there are a few other things that are the result of my Master Assignment, such as some open-source software. In this chapter each of these deliverables will be described.

2.1 RIPE69 Meeting Presentation

The results of the study were accepted to be presented during the RIPE69 Meeting in London. RIPE[1] is the maintainer of the measurement platform that was used to do the measurements for this study, named RIPE Atlas [2]. This made the meeting interesting and relevant. Further details about this meeting and the slides can be found in Appendix A.

2.2 TMA 2015 Paper

The paper that is to be submitted to the The 7th International Workshop on Traffic Monitoring and Analysis or TMA 2015[3]. This paper is the core of this Master Thesis and is included in Appendix B.

2.3 Open-source software and data

During the course of my Master Assignment I have written a software library and two tools to assist in the analysis of the measurement data. These tools were released as open-source so that they can be used or improved by others. In the rest of this section these tools will be described.http://frontpage.fok.nl/nieuws/678308/1/1/50 geschokt-door-reorganisatie-ing.html

2.3.1 BView Parser

This library is written to parse data that is formatted according to the MRT Routing Information Export Format as specified in RFC6396. The Remote Route Collectors (RRCs) operated by RIPE provide detailed routing information about the Internet in this format.

The goal of this library is to make it easy to read a BView file and then access the data that it contains in an object-oriented way. While there are other libraries available, they contain a lot of legacy code and are written in a language other than Java, making them less easy to run on multiple platforms than the library that I wrote. The library is far from complete but it has progressed far enough to allow the extraction of the necessary data for IP2ASN conversion. The code is open-source and available at `https://bitbucket.org/woutifier/bviewparser`

2.3.2 IPASNExporter

Using the BView Parser library this standalone tool, named IPASNExporter, can automatically download the BView data published by RIPE and convert it into a CSV-file containing IP/ASN pairs. An example of the output is shown in Table 2.1. The output can be used to link almost any IP-address to it's corresponding Autonomous System (AS). The code is open-source and available at `https://bitbucket.org/woutifier/ipasnexporter`.

Table 2.1: Example output of IPASNExporter

IP	MASK	ASN	COUNT
82.7.16.0	20	5089	3
112.105.48.0	20	18049	1
64.81.198.0	24	18566	92
60.52.49.0	24	4788	96

[1] https://www.ripe.net/

[2] https://atlas.ripe.net/

[3] http://tma-2015.cba.upc.edu/

2.3.3 Measurement data

The data that was gathered during this Master Assignment is made available publicly through the RIPE Atlas[4] project. Anyone with an interest in either validating my results or in doing his or her own measurements is able to do so. The measurements can be retrieved using the RIPE Atlas API using the measurement IDs provided at `https://github.com/Woutifier/RipeData`

[4] https://atlas.ripe.net/

Chapter 3

Reflection

The final part of one's period as a Master student is one in which a student has a lot of responsibility. This responsibility comes with advantages, such as learning a lot, and disadvantages, such as your inability to properly plan something. This chapter will briefly reflect on some of the lessons that I have learned during the course of my Master Assignment.

Motivation

The Master Assignment normally runs over a period of roughly 6 months, however, this is based on working on it full time. In my case I worked partially at a company and partially at the university. As a consequence instead of 6 months it took me 9 months to finish. 9 months felt like a long time to work on single subject and I have learned that it is necessary to switch the aspect of the subject that you work on regularly in order to stay motivated. One important lesson is that the ability to decide, by yourself, what tools you use and what you work on, on a day to day basis, is of great value.

Planning

Planning your work as a student is something that you never really appear to master, atleast it seems this way to me. I have made numerious schedules that I abandoned after a while. On one hand this was caused by the changing of the direction of the study, which forced a change in the schedule. On the other it was difficult to stick to deadlines especially due to the lack of consequences of exceeding one. The deadline for my graduation, which was caused by the fact that my supervisor had limited time to attend it, was a great help.

Speaking

Perhaps the most challenging aspect of my assignment was the point where I had to present my research in front of more than 500 people, in London. It must have been one of the most nervous moments of my life when my name was announced by the chair. It is very good to know that I have it in me to present, comprehensibly, for such a large audience.

References

[1] ANP/NU.nl: Ook ABN Amro getroffen door DDoS-aanval. http://www.nu.nl/internet/3401202/abn-amro-getroffen-ddos-aanval.html, accessed on 8 Februari 2014

[2] Cowan, C., Wagle, F., Beattie, S., Walpole, J.: Buffer overflows: attacks and defenses for the vulnerability of the decade. In: Proceedings DARPA Information Survivability Conference and Exposition. DISCEX'00. vol. 2, pp. 119–129. IEEE Comput. Soc (2000), http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=821514

[3] Douligeris, C., Mitrokotsa, A.: DDoS attacks and defense mechanisms: classification and state-of-the-art. Computer Networks 44(5), 643–666 (Apr 2004), http://www.sciencedirect.com/science/article/pii/S1389128603004250

[4] Ferguson, P., Senie, D.: Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing. RFC 2827 (Best Current Practice) (2000), http://www.ietf.org/rfc/rfc2827.txt

[5] He, Y., Faloutsos, M., Krishnamurthy, S., Huffaker, B.: On routing asymmetry in the Internet. In: Global Telecommunications Conference, 2005. GLOBECOM '05. IEEE. vol. 2, pp. 6 pp.– (Nov 2005)

[6] He, Y., Faloutsos, M., Krishnamurthy, S.V.: Quantifying routing asymmetry in the Internet at the AS level. In: GLOBECOM. pp. 1474–1479. IEEE (2004), http://dblp.uni-trier.de/db/conf/globecom/globecom2004.html#HeFK04

[7] Hoque, N., Monowar, B., Baishya, R., Bhattacharyya, D., Kalita, J.: Network attacks: Taxonomy, tools and systems. Journal of Network and Computer Applications (2013)

[8] Infosecurity: Dutch banking system DDoS'd. http://www.infosecurity-magazine.com/view/31663/dutch-banking-system-ddosd/, accessed on 8 Februari 2014

[9] Karig, D., Lee, R.: Remote Denial of Service Attacks and Countermeasures. Princeton University Department of Electrical Engineering Technical Report CE-L2001-002 (October) (2001)

[10] Li, J., Mirkovic, J., Wang, M., Reiher, M., Zhang, L.: SAVE: Source address validity enforcement protocol. In: Proceedings - IEEE INFOCOM. vol. 3, pp. 1557–1566 (2002)

[11] Lippmann, R., Fried, D., Graf, I., Haines, J., Kendall, K., McClung, D., Weber, D., Webster, S., Wyschogrod, D., Cunningham, R., Zissman, M.: Evaluating intrusion detection systems: the 1998 DARPA off-line intrusion detection evaluation. In: Proceedings DARPA Information Survivability Conference and Exposition. DISCEX'00. vol. 2, pp. 12–26. IEEE Comput. Soc (2000), http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=821506

[12] Mirkovic, J., Reiher, P.: A taxonomy of DDoS attack and DDoS defense mechanisms (2004)

[13] NOS: Probleem banken door cyberaanval. http://nos.nl/artikel/492603-geslaagde-cyberaanval-op-banken.html, accessed on 8 Februari 2014

[14] Park, K., Lee, H.: On the effectiveness of route-based packet filtering for distributed DoS attack prevention in power-law internets (2001)

[15] Peng, T., Leckie, C., Ramamohanarao, K.: Survey of network-based defense mechanisms countering the DoS and DDoS problems. ACM Computing Surveys 39(1), 3–es (Apr 2007), http://dl.acm.org/citation.cfm?id=1216370.1216373

[16] Prince, M.: Technical Details Behind a 400Gbps NTP Amplification DDoS Attack. http://blog.cloudflare.com/technical-details-behind-a-400gbps-ntp-amplification-ddos-attack, accessed on 19 Februari 2014

[17] Prince, M.: The DDoS That Almost Broke the Internet. http://blog.cloudflare.com/the-ddos-that-almost-broke-the-internet, accessed on 27 January 2014

[18] Prolexic: Quarterly Global DDoS Attack Report Q3. http://www.prolexic.com/knowledge-center-ddos-attack-report-2013-q3.html, accessed on 27 January 2014

[19] Ristic, I.: Programming Model Attacks. In: Dageforde, M. (ed.) Apache Security, chap. 5.4.3. O'Reilly Media, Inc. (2005)

[20] RTL-Nieuws: Rabobank getroffen door DDoS-aanval. http://www.rtlnieuws.nl/nieuws/binnenland/rabobank-getroffen-door-ddos-aanval, accessed on 8 Februari 2014

[21] Schuba, C., Krsul, I., Kuhn, M., Spafford, E., Sundaram, A., Zamboni, D.: Analysis of a denial of service attack on TCP. In: Proceedings. 1997 IEEE Symposium on Security and Privacy (Cat. No.97CB36097). pp. 208–223. IEEE Comput. Soc. Press (1997), http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=601338

[22] Schwartz, Y., Shavitt, Y., Weinsberg, U.: On the Diversity, Stability and Symmetry of End-to-End Internet Routes. In: INFOCOM IEEE Conference on Computer Communications Workshops , 2010. pp. 1–6 (2010)

[23] University of Twente: Master guide 2012/2013, http://www.utwente.nl/csc/programmeinformation/rules_documents/guides/studiegids2012_msc_csc-hmi-mte.pdf

[24] Zhang, B., Liu, R., Massey, D., Zhang, L.: Collecting the internet AS-level topology (2005)

APPENDIX A

RIPE69 MEETING PRESENTATION

RIPE69 is a meeting for Internet Service Providers (ISPs), network operators and other interested parties (such as academics). It is a five-day event with presentations covering a wide range of network related topics. The event was held in the Novotel in London West (UK). My research was presented in the plenary session (recording is available) and the slides that were used are included in this Appendix.

Name:	RIPE69 Meeting
Deadline:	1st of October 2014
Date:	3rd until 7th of November 2014
Venue:	Novotel London West Hotel, London, United Kingdom
Video recording:	`https://ripe69.ripe.net/archives/video/163/`
General Website:	`https://ripe69.ripe.net`

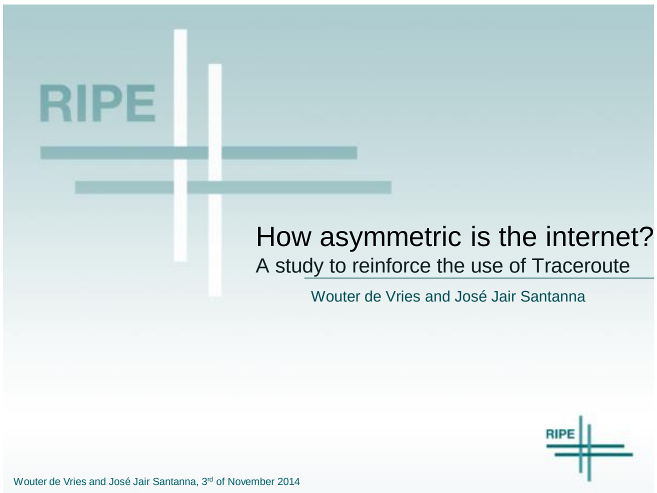

How asymmetric is the internet?

A study to reinforce the use of Traceroute

Wouter de Vries and José Jair Santanna

Wouter de Vries and José Jair Santanna, 3rd of November 2014

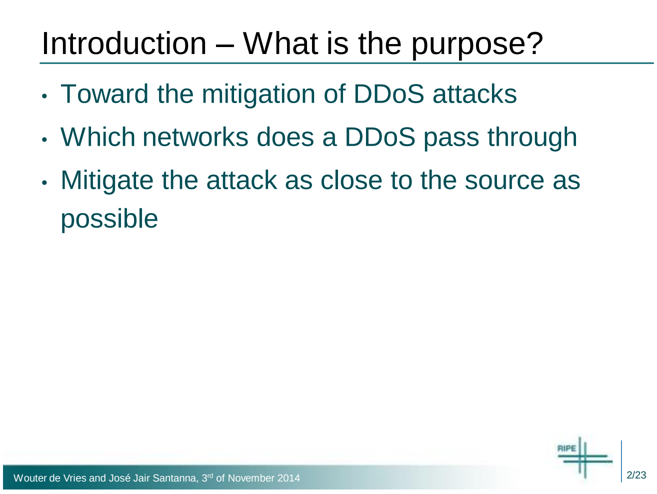

Introduction – What is the purpose?

- Toward the mitigation of DDoS attacks
- Which networks does a DDoS pass through
- Mitigate the attack as close to the source as possible

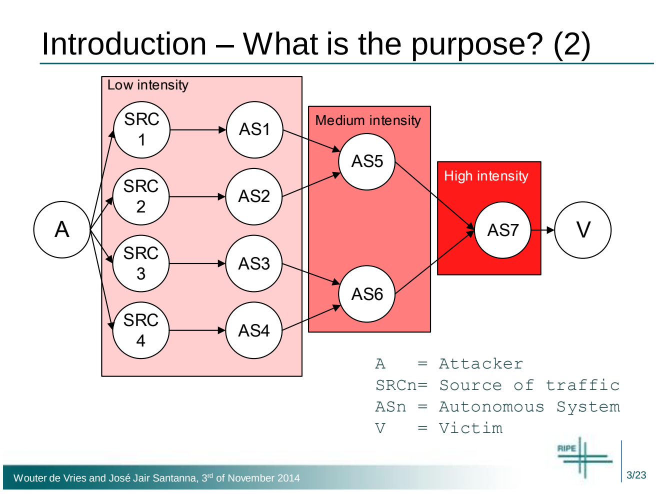

Introduction – What is the purpose? (2)

```
A    = Attacker
SRCn= Source of traffic
ASn  = Autonomous System
V    = Victim
```

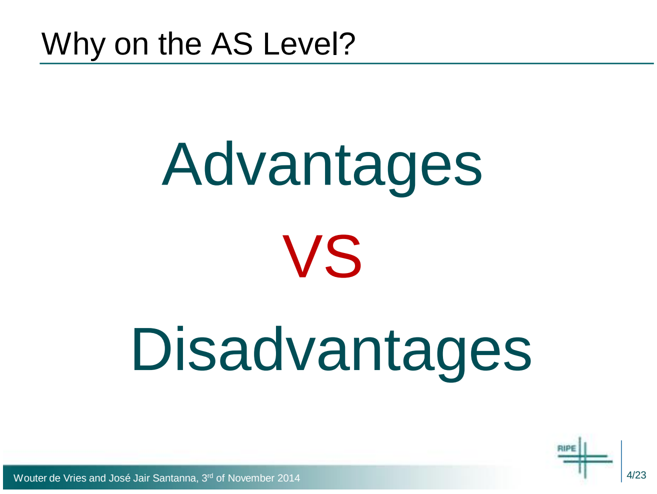

Why on the AS Level?

Advantages

VS

Disadvantages

Introduction - Problem

- Measuring the reverse path of an attacker is less trivial than it seems
- Path can be, and often is, asymmetric[1][2][3]

[1] Y. He, M. Faloutsos, S. Krishnamurthy, and B. Huffaker, “On routing asymmetry in the Internet,” in Global Telecommunications Conference, 2005. GLOBECOM ’05. IEEE, vol. 2, Nov. 2005
[2] Y. He, M. Faloutsos, and S. V. Krishnamurthy, “Quantifying routing asymmetry in the Internet at the AS level.” in GLOBECOM. IEEE, 2004
[3] Y. Schwartz, Y. Shavitt, and U. Weinsberg, “On the Diversity, Stability and Symmetry of End-to-End Internet Routes,” in INFOCOM IEEE Conference on Computer Communications Workshops, 2010

Introduction – Research question

How asymmetric is the internet?

Expectation:

The internet is more symmetric near the source and target

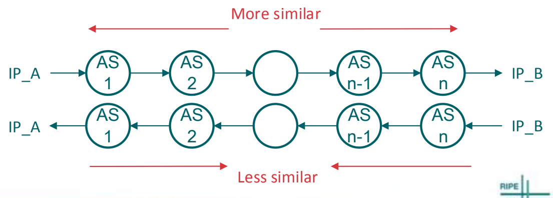

Methodology

- Large scale measurements, using RIPE Atlas
- 4000 out of 7000* probes located worldwide
- 2000 pairs of probes to measure paths between
- Thanks to RIPE Atlas our UDM limit was greatly increased

*Approximately

Selected probe distribution

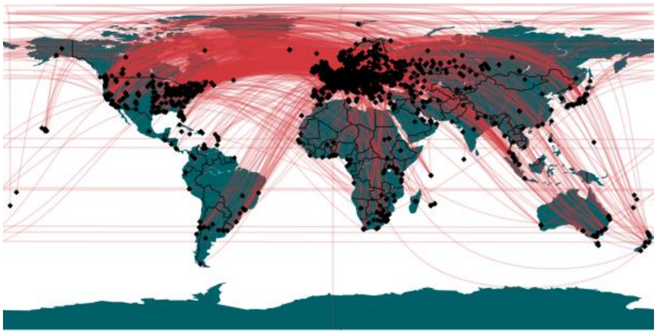

Pairs “randomly” selected

Methodology – Pairs

Continent	Probe count	Fraction
Europe	2.681	67,03%
North America	724	18,10%
Asia	267	6,68%
Africa	157	3,93%
Oceania	109	2,73%
Others	62	1,55%
Total	**4.000**	**100%**

Experiment setup

- Traceroute
- Two measurements every three hours per pair
- Ten days
- 80 samples per pair
- 160.000 measurements
- 5.256.138 records
- ~1 gigabyte

Analysis – Determining ASN

- BGP Routing tables (Provided by RIS/RIPE)
- Process to list of IP-range/ASN tuples
- Use binary tree to quickly match IP address to ASN with longest prefix matching

Analysis – Determining ASN (2)

- Tool to process binary BGP routing table
  http://bit.ly/XXRsxd
- Tool to build a list of IP-range/ASN tuples
  http://bit.ly/1tJDoTt

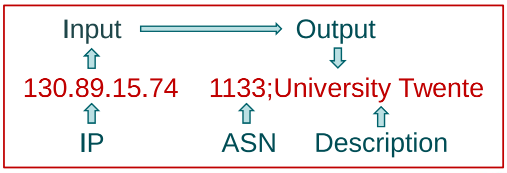

Results – Path length on the AS level

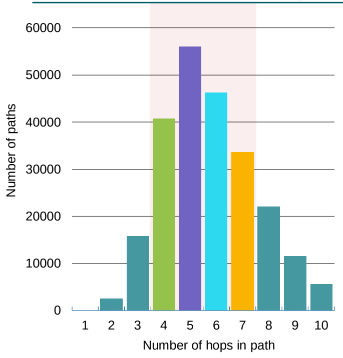

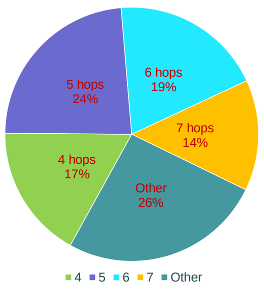

Results – The numbers

- 119.550 measured network paths
- 2.275 unique Autonomous Systems in total
- 1.717 contain probes
- 12.6% (15.053) of pairs are symmetric

Results – Change of path over time

- Compare paths from A to B over time
- Determine Levenshtein distance to make two paths equal
- Compare each path to first path

For example path from A to B with 1 insert/delete/change operation

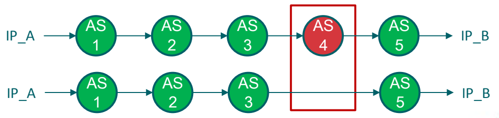

Results – Change of path over time (2)

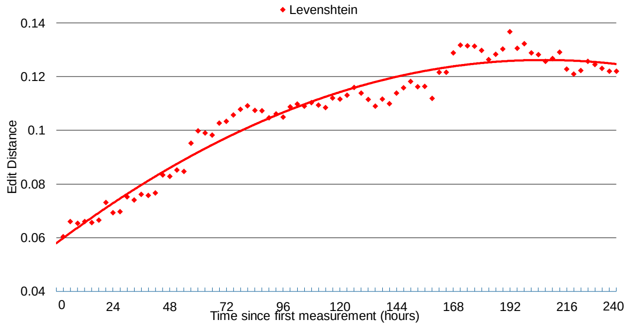

Network paths barely change over time*

Results – Equal Consecutive Hops

- Count number of Equal Consecutive Hops (ECH) from left to right and right to left.
- Group by path length
- First and last hop are always equal

2 + 1 = 3 Consecutive equal hops

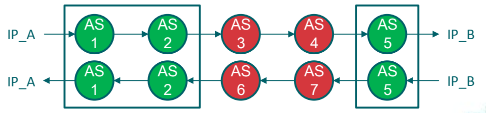

Results – Equal Consecutive Hops (2)

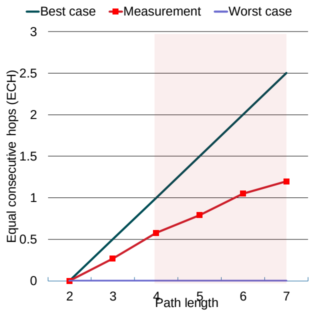

#	Samples	Std. Dev.	ECH
2	792	0	0
3	8164	0.498	0,27
4	20.038	0,733	0,58
5	27.870	1,013	0,79
6	23.442	1,299	1,05
7	17.223	1,492	1,2

One extra chance to mitigate DDoS attacks

Results – Equality by position

- Compare hop in forward path to the one in the same position in reverse path
- Either equal (1) or not equal (0)
- Average per position

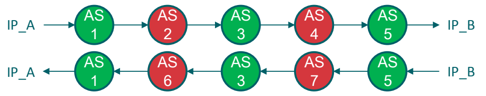

Results – Equality by position (2)

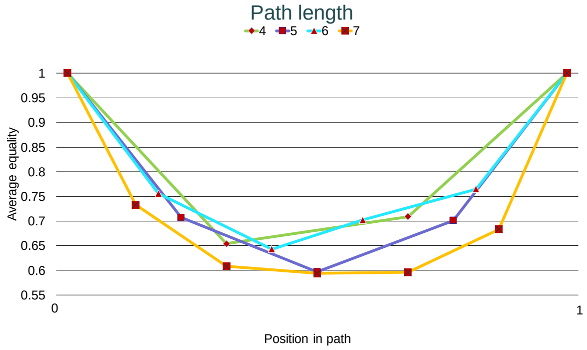

High chance to have equal hops

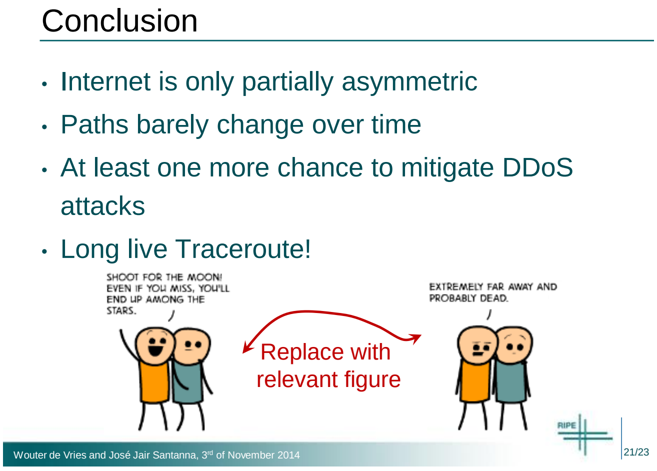

Conclusion

- Internet is only partially asymmetric
- Paths barely change over time
- At least one more chance to mitigate DDoS attacks
- Long live Traceroute!

SHOOT FOR THE MOON! EVEN IF YOU MISS, YOU'LL END UP AMONG THE STARS.
EXTREMELY FAR AWAY AND PROBABLY DEAD.

Replace with relevant figure

Wouter de Vries and José Jair Santanna, 3rd of November 2014

21/23

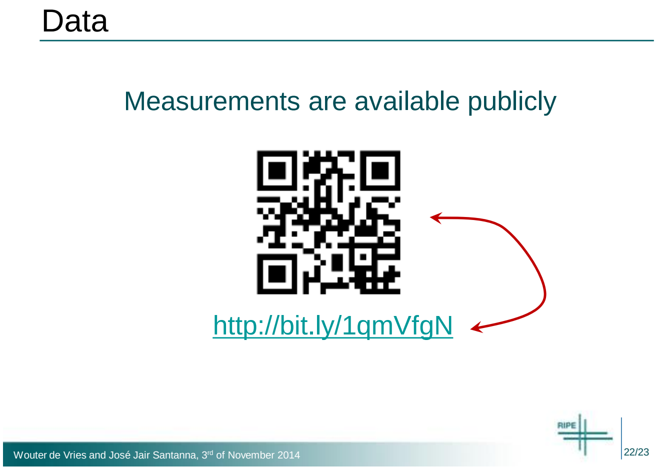

Data

Measurements are available publicly

http://bit.ly/1qmVfgN

Wouter de Vries and José Jair Santanna, 3rd of November 2014

22/23

Questions?

Comments are also allowed
(despite the large question mark)

Appendix B

TMA 2015 Paper

The 7th International Workshop on Traffic Monitoring and Analysis or TMA 2015 is a conference on virtually all types of (Internet) Network Measurements. The conference will be held on Campus Nord of the Universitat Politècnica de Catalunya in Barcelona, Spain. The paper that will be submitted to this conference is the main product of the graduation phase of my Master's degree and it is included in this Appendix.

Name:	TMA 2015 - International Workshop on Traffic Monitoring and Analysis
Deadline:	28th of November 2014
Date:	23rd until 24th of April 2015
Venue:	Campus Nord, Universitat Politècnica de Catalunya, Barcelona, Spain
General Website:	`http://tma-2015.cba.upc.edu/`

How asymmetric is the Internet?
A Study to Support DDoS Mitigation Approaches

Wouter de Vries, José Jair Santanna, Anna Sperotto, and Aiko Pras

University of Twente
Design and Analysis of Communication Systems (DACS)
Enschede, The Netherlands
{w.b.devries-1@student.utwente.nl;j.j.santanna@utwente.nl
a.sperotto@utwente.nl;a.pras@utwente.nl}

Abstract. A promising approach to mitigate large Distributed Denial-of-Service (DDoS) attacks is to mitigate them closer to the source. To do this it is necessary to determine the network paths that the attackers use. A network path is a path that a packet takes to reach its target. However, determining the network path that an attacker used to reach its target is less trivial than it appears. Tools such as Traceroute allow the user to determine the path towards a target (i.e. the forward path), but not the path from the target to the source (i.e. the reverse path) due to routing asymmetry. Routing asymmetry means that the network path between two hosts may be different in opposite directions. Although previous studies have shown that this asymmetry is widespread, a more detailed characterization is lacking. In this paper routing asymmetry is investigated in depth using world wide large scale measurements using 4.000 probes. The main goal of this paper is to provide characteristics about Internet asymmetry, with the possible application of DDoS mitigation. However, our findings contribute to a conclusive overview of Internet asymmetry, which assist researchers and engineers in making valid assumptions about the asymmetry of network paths.

Keywords: Internet, Measurements, Asymmetry, Ripe Atlas

1 Introduction

Distributed Denial-of-Service (DDoS) attacks are a constant threat on the Internet. Existing mitigation solutions have limitations against bandwidth intensive attacks (i.e. volumetric attacks), especially when faced with attacks at scale of hundreds of gigabits per second. A volumetric attack can overwhelm the network leading to the victim before the mitigation solution can be effective. Attacks of this type and on this scale have recently taken place on CloudFlare with more than 300 Gbit/s [11, 12]. A promising approach for mitigating attacks on this scale is blocking them closer to the sources that generate the traffic [9, 14]. This approach has the advantage of preventing a single network from having to absorb all the DDoS traffic.

Although this mitigation approach can be effective, it requires a method to reliably determine the networks that an attack passes through (i.e. the network path), from the viewpoint of the attack target. This could be done by measuring the path from the target to the attacker (i.e. the reverse path) using a tool such as Traceroute and assuming that this is the same path that the attacker took to reach the victim (i.e. the forward path). However, this method has proven to be inaccurate because of routing asymmetry [4, 5, 10, 13]. This means that the forward path may not be the same as the reverse path. However, existing studies lack the analysis of where asymmetry occurs.

In this paper we look into the asymmetry of network paths. We investigate to what extent the forward path can still be determined using the reverse path if the characteristics of Internet asymmetry are known. The goal of this study is to analyse Internet routing asymmetry in depth, with the application of DDoS mitigation in mind. To perform this analysis we measure network paths between 4000 probes across the world. We analyze the resulting data, in depth, for network path asymmetry from the Autonomous System (AS) level. We show that most routes are not completely symmetrical but that they do have properties that make them useful for specific applications, such as DDoS mitigation. The contribution of this paper is providing information that researchers and engineers can use for the practical applicability of forward/reverse paths.

This paper is organized as follows. In Section 2 we discuss the related work followed by Section 3, where we explain our hypothesis and requirements. In section 4 we describe our methodology. Then, in Section 5, the analysis will be described. Finally, we will present our conclusions in Section 6.

2 Related work

Researchers have been studying Internet routing asymmetry for some time [5] [4] [13] [7]. In this section we will discuss a few studies that have investigated the level of routing asymmetry on the Internet and indicate what shortcomings they have that we aimed to solve.

First, the research in [5] on route asymmetry covers the AS level. They conclude that route asymmetry, on the AS level, is only present in approximately 14% of the routes. However, this research is based on results gathered using the Active Measurement Project (AMP) which runs mainly on academic networks and uses only 135 probes. In their follow up study [4], they use 350 probes selected from 1200 public traceroute servers. They note that the routing asymmetry percentage is much higher on commercial networks, namely 65%, which negatively impacts the usability of Traceroute to measure reverse network paths.

In addition, while they have conducted extensive research on route asymmetry on the AS level they have not looked at the relative position of asymmetry (e.g. close to the target of the traceroute, in the middle or close to the source of the traceroute). If we are interested in the remaining usability of reverse paths this is an interesting measurement, for example for applications that do not require the entire path to be symmetric. They proposed an interesting framework for quantifying the change in paths in which they use the the Levensthein Edit Distance (ED) algorithm as a way to determine the distance between two paths.

Secondly, research in [13] concluded that the asymmetry on the AS level is substantially higher than in [4] [5]. According to them, asymmetry on the AS level is as high as 90%. The cause of this difference could, for example, be that this study was conducted 5 years later or that their dataset is obtained using only a total of 220 probes.

Finally, the research in [7] proposes a way of determining the actual path that a packet has taken to reach a point in a network, with routing asymmetry in mind, from the viewpoint of the receiver. They do this mainly for troubleshooting purposes (e.g. which network is dropping packets). Their method involves a system of widely deployed probes, IP spoofing and the use of the record route option in the IP header. While the theory behind this method is sound, it is difficult to deploy in practice for a few reasons. First, potential users need to have widely deployed probes in place. Secondly, their method uses the Record Route option in the IP header. However, this option is widely ignored [6] and packets that use this option are often dropped. Finally, the use of IP spoofing, the act of replacing the source address in a packet with a forged one, can be problematic due to issues with company policies, ethics and the fact that there are techniques to block IP spoofing such as proposed in Request for Comment (RFC) 2827 [2], which is currently known as Best Current Practice (BCP) 38.

To summarize, existing studies have conflicting results, but all of them agree that asymmetry is a common occurence. The difference in the results may be caused by the limited number of probes used or the type of network (e.g. academic vs commercial). Furthermore their analysis of routing asymmetry is limited, for example, they do not determine in which parts of the path asymmetry occurs most often.

3 Hypothesis and Requirements

The goal of the first part of this section is to describe some terminology to introduce our hypothesis. Next, we will describe the requirements for the measurements needed to answer it.

In this paper we consider a *network path* an ordered list of networks that connect two end-systems on the Internet. Although there are studies that differentiate networks by IP address or even as IP address range [4], we chose to represent networks as Autonomous System (Autonomous System (AS)). ASes are easier to be linked to an owner that can be contacted to help mitigate DDoS attacks and it is trivial to cluster IP addresses that belong to the same administrative network.

As shown in Fig. 1 there are two distinct paths between a pair of end-systems A and B: The forward path and the reverse path. When both paths are completely equal then the path is symmetric, if it is not equal then it is asymmetric. Note that in order for Traceroute to be reliable in determining a complete network path from the viewpoint of the receiver, the Internet would have to be completely symmetric. Our goal is to show that network paths are symmetric near the end-systems to show reliable locations for mitigation. We defined the following hypothesis: DDoS attacks can be mitigated outside the victims network using network paths measured using Traceroute.

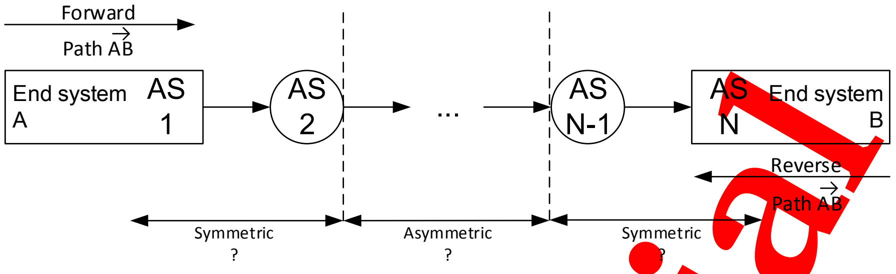

Fig. 1: Network path

3.1 Requirements

The main requirement for investigating our hypothesis was having a large amount of Internet connected computer systems which we could control.

In order to meet our main requirement we use RIPE Atlas. This project manages probes around the world for the specific purpose of network measurements. A probe is a dedicated network measurement device that can be placed in a network to allow measurements to be performed remotely. The Atlas project consists of approximately 7.000 distributed probes[1] worldwide. Although we are aware of several other measurement infrastructures, such as PlanetLAB[2], EmanicsLAB[3] and the NLNOG Ring[4], these do not provide the scale and distribution that was required for measurements that are representative of the Internet.

RIPE Atlas has imposed a credit system that limits measurements in three ways. The credits that are consumed per day, the number of measurements that can be run concurrently and the total number of credits that can be consumed. These limits have a consequence on the number of probes that can be used and in which combination. Credits can, for example, be earned by hosting a RIPE Atlas probe.

Due to the credit limit not all probes that are available can be used. This further depends on the measurement layout, which probe measures what and to what other probe.

[1] RIPE Atlas System Statistics: https://atlas.ripe.net/
[2] https://www.planet-lab.org/
[3] http://www.emanicslab.org/
[4] https://ring.nlnog.net/

4 Methodology

In this section we will describe the methodology that we used. We will first explain the way in which the probes are selected and then how they are configured.

We considered three layouts in which the probes can conduct the measurements. Note that to be able to determine route asymmetry between two probes, each probe has to traceroute the other. In the considered layouts each probe performs traceroutes to the probes to which it is connected.

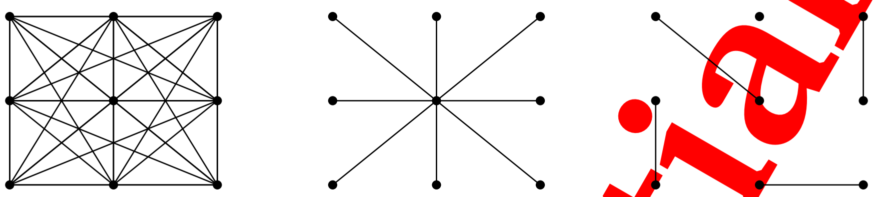

(a) Fully connected

(b) Star

(c) Random Tuples

Fig. 2: Probe layout

Fully connected layout (Fig. 2a) - This layout has the advantage of utilizing the complete potential of the involved probes, every probe measures the path to every other probe. The disadvantage is that due to the credit limit only a very limited amount of probes[5] from the total can be used. A small amount of probes means that specific network issues that occur at individual probes have a large impact.

Star layout (Fig. 2b) - In comparison to the fully connected topology this has the advantage of allowing many more probes to be used. However, in this case the center probe will have a large impact on each measurement. Network issues at the center probe can cause the entire experiment to fail.

Random tuple layout (Fig. 2c) - In this layout random tuples of probes are selected. This has the advantage of minimizing the impact of a single misbehaving probe. Furthermore, it allows for a much larger selection of probes, considering the Atlas limits. Because of these advantages this is the combination method that we used.

Using the random tuple layout we selected 4.000 probes, meaning 2.000 tuples, in a way that favoured longer geographic distances. The attempt to have longer geographic distances is to prevent a large concentration of probes in Europe, as most probes are located there. The algorithm used to select the probes works by randomly picking probes and comparing the distance between them to some threshold (10.000), if the threshold is exceeded then the probe tuple is added to the final result set. If, after a number of attempts (2.000), no probe

[5] 112 probes considering the 1 million credit limit

tuples can be found that exceed the threshold then the threshold will be lowered. This algorithm in pseudocode is shown below.

```
function selectrandomtuples($\text{probeset}_{\text{available}}$, threshold):
    Variables:probetuples = []
    WHILE len(probetuples) <2000:
        probe1 = randomfrom($\text{probeset}_{\text{available}}$),probe2 = randomfrom($\text{probeset}_{\text{available}}$)
        attempts = 0
        WHILE distance(probe1, probe2) <threshold OR attempts >= 2000:
            probe2 = randomfrom($\text{probeset}_{\text{available}}$)
            attempts += 1
        IF attempts >= 2000:
            threshold -= 10
            CONTINUE
        probetuples += tuple(probe1, probe2)
        $\text{probeset}_{\text{available}}$ -= probe1
        $\text{probeset}_{\text{available}}$ -= probe2
    RETURN probetuples
```

The distribution over continents in terms of numbers is shown in Table 1. There is a large skew towards Europe which is caused by the relatively large number of probes located there. The average distance between two probes in a tuple is 6.945 kilometres (as the crow flies).

Table 1: Distribution over continents

Continent	**Selected**	**Available**	**Fraction**	**Fraction of selected**
Europe	2681	5200	51.56%	67.03%
North America	724	1003	72.18%	18.10%
Asia	267	420	63.57%	6.68%
Africa	157	223	70.40%	3.93%
Oceania	109	145	75.17%	2.73%
South America	59	87	67.82%	1.48%
Antarctica	1	1	100.00%	0.03%
Unknown	*2*	*4*	*50.00%*	*0.05%*
Total	***4000***	***7083***		***100%***

For every selected pair consisting of probe A and probe B two measurements were scheduled. One measurement, consisting of a traceroute, was configured from probe A to probe B (the forward path) and another from probe B to probe A (the reverse path). Note that in the ideal case, when the path is completely symmetric, the forward path is the exact opposite of the reverse path.

Network variances over time were smoothed out by scheduling the measurements to run every three hours, for ten days. This was limited by the total amount of credits we were allowed to consume. The measurements were performed from 14:00 on the 28th of July 2014 to 14:00 on the 7th of August 2014, Coordinated Universal Time (UTC).

Our choice of probes was optimized to prevent a large cluster of probes in Europe by increasing the geographic distance between pairs. However, this relation may have caused a bias in network path length. In order to show that this is not the case we plot the geographic distance against the number of hops in the forward network path on the AS level. The geographic distance is calculated as the great circle distance (i.e. as the crow flies)[6]. The result of this is shown in Fig. 3. As we expected there is no clear relation between the geographic distance between probes and the number of hops in the path between them. This led us to the conclusion that there is no significant relation between geographic distance and network path length.

In Fig. 4 the distribution of the length of the measured paths is shown. The length of the forward path and the length of the reverse path are counted individually. Most paths contain five different AS-numbers. This indicates that in those cases three autonomous systems aside from the one the receiver and the sender are in (e.g. their ISPs) are involved in routing the packets.

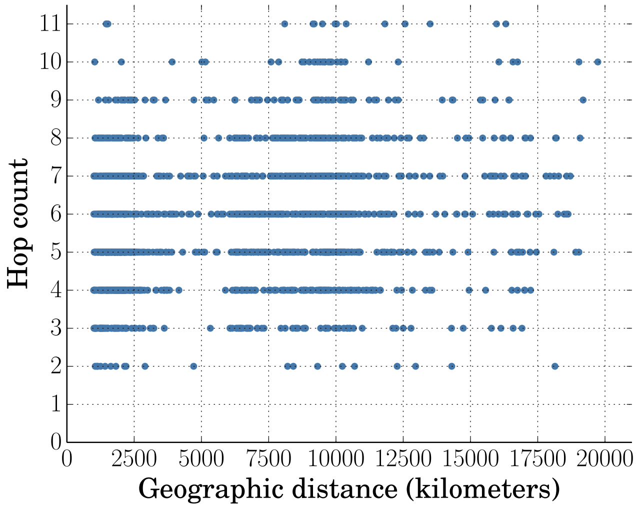

Fig. 3: Geographic distance vs path length

[6] The great circle distance is the shortest distance between two points on a sphere

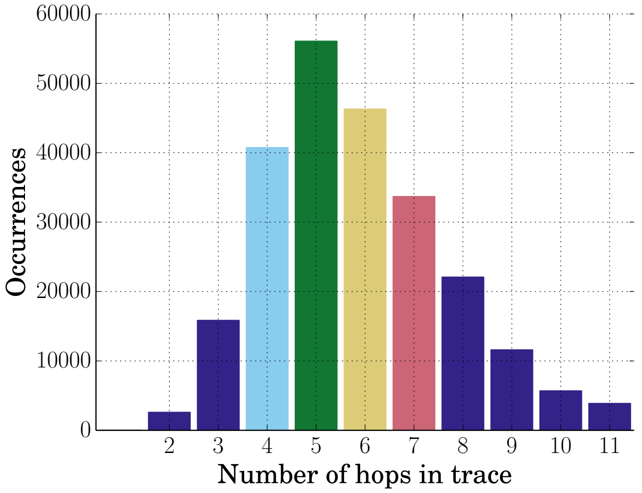

Fig. 4: Distribution of path length

The measurements that were performed by the probes were not completely perfect or complete. This may, for example, have been caused by probes that delayed their measurements for too long or did not perform them at all. Because of this we have had to apply some filters to the data set. Prior to the filtering we had 153,638 potential forward/reverse path pairs, of a theoretical 160,000. The absence of some paths can be attributed to some probes that did not respond.

Since forward and reverse path measurements are initiated from two different probes there can be a delay between the two measurements being executed. To prevent this difference from influencing the results we only match paths if they were measured no more than 600 seconds (10 minutes) apart. If they are outside that time limit the forward/reverse pair is discarded. This prevents path instability from being interpreted as path asymmetry. This filter reduced our potential forward/reverse path pairs by 16,103.

The second filter we implemented is based on the principle that the first hop of the forward path should be the last hop of the reverse path, because these are the origin and destination networks. The same principle applies in the opposite direction. Measurements where this is not the case can be caused by incomplete traces. We filtered all forward/reverse path pairs where this was the case. This removed 14,620 results from the set.

To prevent probes that measured completely empty paths to influence the results we filtered all pairs that contained a completely empty path. Completely empty paths do not exist in actual networks, as a network path always contains at least a single hop, even if the source and target IP addresses are in the same network. This can be caused by incomplete traces or probes that are not executing their measurements. This filter reduced our result set by 3,365.

The three filters that we implemented removed a total of 34,088 results from our data set. Leaving 119,550 or 74.72% of the theoretical 160,000 pairs.

For paths that contained unresolvable hops we considered a few options. The first option is to discard all path pairs that contained such a hop. However, this would impact a significant part of the result set as unresolvable hops are common. Another option, which was also implemented in [3] is to simply consider an unresolvable hop as a wild card, meaning that it will match any hop in the opposite path that is in the same position.

RIPE Atlas probes conduct their traces on the IP level where each hop consists of a single IP address. Because we want to look at the network paths from the AS level it was necessary to convert the measured paths. In order to convert IP addresses in a fast way we used the BGP routing table dumps obtained from the Remote Route Collector (RRC)s managed by the Routing Information Service (RIS), which in turn is operated by RIPE. These routing tables contain a large amount of routes that are announced on the Internet by different ASes. Using these routes we are able to determine the AS number for a given IP range. The tool we used for this and its source is available online[7].

Each IP address in the paths on the router level was converted to their corresponding AS number. It is apparently common for multiple hops to occur within the same network. This is shown by the reduction in the number of hops in network paths on the router level in comparison to network paths on the AS level, which is, on average, 64.46%.

5 Analyses

In this section we analyse the dataset that was obtained using the methodology described in the previous section. Our dataset contains a total of 2275 unique AS numbers, of which 1717 contain one or more probes. Of all results in our dataset, 15053 (12.6%) forward/reverse path tuples are completely symmetric and 104497 (87.4%) show asymmetry. This is similar to the results found in [13], however, we use far more probes. The large percentage of asymmetric paths amplifies the importance of understanding the characteristics of Internet asymmetry.

Before we start the analysis we introduce two metrics for calculating the Edit Distance (ED) between two paths. One is the Levenshtein algorithm [8] which was first used for this purpose in [4] [5]. The Levenshtein algorithm counts the number of required insert, delete or change operations to make two paths equal to each other. The Levenshtein algorithm was originally intended to be used to measure the differences between strings, however, it can be used without modification for measuring the change in network paths. In addition to the Levenshtein algorithm we also use a variation called Damerau-Levenshtein [1]. Damerau-Levenshtein extends the original algorithm by also counting transpose operations as a single change. It is much less sensitive to swapped hops. The extended algorithm is interesting in contexts where the presence of ASes on a path are of more importance than their specific location.

[7] IPASNExporter: https://bitbucket.org/woutifier/ipasnexporter

5.1 Stability over time

We begin our investigation by determining the change of paths over time. This is of interest because it is not always possible to measure the reverse path at the exact time that the forward path was established. We calculate the average ED over all paths over time. The value on the Y-axis, the ED, is determined as follows: The first path to a destination is taken as a ground truth to which each consecutive path will be compared. We then calculate the ED based on the Levenshtein algorithm as shown in Fig. 5. We had to modify the algorithm slightly because not all paths are of the same length, which would cause longer paths to have a much higher impact on the results than shorter paths. Therefore, we normalize the ED by dividing it by the path length as shown in formula 1.

$$\frac{ED(forward, reverse)}{MAX(len(forward), len(reverse))} \tag{1}$$

The normalized ED is between 0.0 (i.e. completely symmetric) and 1.0 (i.e. completely asymmetric). Note that the graphs show that network paths are not subject to great change over time. We compared the results using the Levenshtein algorithm to the Damerau-Levenshtein algorithm and this showed results which are almost completely identical to Fig. 5, this indicates that the relative position of a network in a path is stable. The instability appears to stop increasing after 8 days, therefore measurements should be done over a longer period of time to show if this behaviour persists.

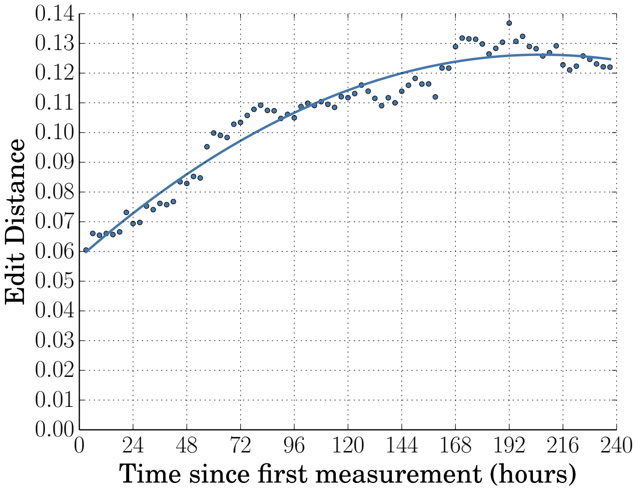

Fig. 5: ED over time using Levenshtein algorithm

5.2 Absolute difference

We look at the absolute difference between the forward and reverse path pairs to get an understanding of how big the impact of routing asymmetry is. In Fig. 6 the difference between all forward/reverse path pairs is shown using both algorithms. Note that the difference between the results of the two algorithms indicates that it is a common occurence for two hops to be swapped in either the forward or reverse path. Furthermore, most forward/reverse path pairs show a distance of either 1 or 2 from their counterpart.

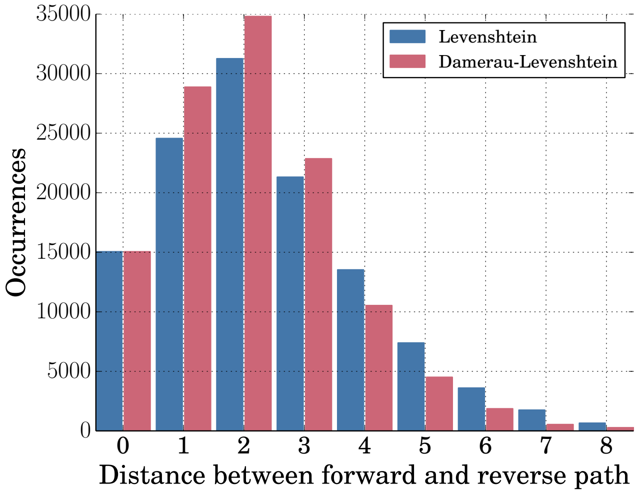

Fig. 6: Distance between forward and reverse path

5.3 Relative difference by position

In this section we show the similarity of hops based on their relative position in the path. This shows if a certain hop is usable for mitigation. If a forward and reverse trace have different lengths then they are not included in this figure, which results in 28139 result pairs being used in Fig. 7. Note that this shows how the symmetry decreases as we move closer to the middle of the path, as expected.

Given this measure of asymmetry we try to find out if the majority of asymmetry is caused by a small number of networks (i.e. ASes). Towards that end we look at which ASes are involved in each point of asymmetry. From the approximately 500 ASes that are involved in some point of asymmetry we see that the top 10 is responsible for 48% of the total. We categorize thse ten ASes in three types: T1 for Tier 1 providers, T2 for Large ISPs and IXP for Internet Exchange Points. The results are shown in Table 3. It is obvious that the largest Internet

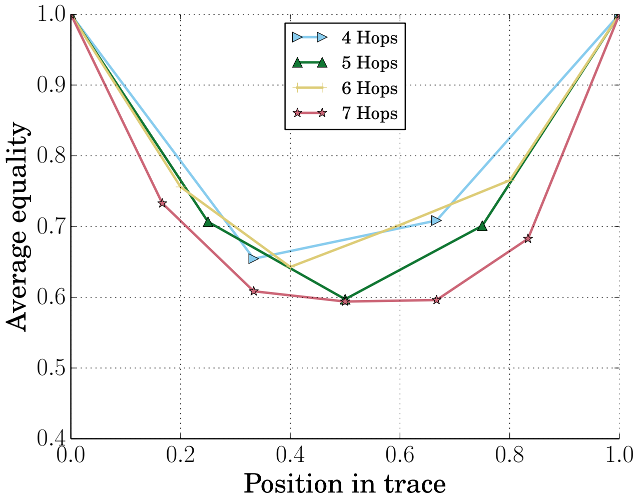

Fig. 7: Average equality by position in trace

Table 2: Average equality by position in trace

#	Average equality
1	1.00
2	1.00 1.00
3	1.00 0.81 1.00
4	1.00 0.65 0.71 1.00
5	1.00 0.71 0.60 0.70 1.00
6	1.00 0.76 0.64 0.70 0.77 1.00
7	1.00 0.73 0.61 0.59 0.60 0.78 1.00

Service Provides (i.e. Tier 1 providers), cause the largest part of the asymmetry. It is likely that this is because those providers are also the ones which have the most peering connections.

5.4 Consecutive equal hops

We calculate the average number of hops from each side of the forward/reverse path that are equal. This approach can be used even if the lengths of the forward/reverse path are unequal. The average number of Consecutive Equal Hops (CEH) is plotted against the number of hops in the forward path in Fig. 8a. The CEH is calculated according to the following algorithm

```
function calculateCEH($\text{path}_{\text{forward}}$, $\text{path}_{\text{reverse}}$):
    Variables:$\text{length}_{\text{total}} = \mathit{LEN}(\text{path}_{\text{forward}})$,$\text{l1} = \mathit{CEIL}(\text{length}_{\text{total}}/2.0)$,
    $\text{l2} = \mathit{FLOOR}(\text{length}_{\text{total}}/2.0)$,$\text{forwardceh} = 0$,$\text{reverseceh} = 0$
    FOR index,hop in $\text{path}_{\text{forward}}$:
```

Table 3: Top 10 ASes involved in asymmetry

Position	ASN	Name	Type
1.	3356	Level 3 Communications, Inc.	T1
2.	174	Cogent Communications	T1
3.	1299	TeliaSonera International Carrier	T1
4.	3257	Tinet SpA	T1
5.	3216	OJSC Vimpelcom	T2
6.	34984	TELLCOM ILETISIM HIZMETLERI A.S.	T2
7.	1200	Amsterdam Internet Exchange B.V.	IXP
8.	2914	NTT America, Inc.	T1
9.	6453	TATA Communications, Inc.	T1
10.	6695	DE-CIX Management GmbH	IXP

```
        IF hop != $\text{path}_{\text{reverse}}$[index]:
            BREAK
        forwardceh++
    reverseinplace($\text{path}_{\text{forward}}$)
    reverseinplace($\text{path}_{\text{reverse}}$)
    FOR index,hop in $\text{path}_{\text{reverse}}$:
        IF hop != $\text{path}_{\text{forward}}$[index]:
            BREAK
        reverseceh++
    RETURN (forwardceh + reverseceh)/2.0 - 1
```

Included in Fig. 8a is the confidence interval. This figure shows that for path lengths 6 and 7 there is on average at least one additional network aside from source and target networks where a DDoS attack can be mitigated. For the most common path length, five, there is another network where an attack can be mitigated in approximately 75% of the cases.

In Fig. 8b only the first complete result for each pair is considered. These graphs show that it is not necessary to do repeated measurements over a longer period of time to determine route asymmetry. Note that this could also indicate that route asymmetry does not vary significantly over time.

6 Conclusion

In this paper we have analysed and characterized several aspects of Internet routing asymmetry. Our analysis has been conducted on a large scale using RIPE Atlas. The results from our study contribute to assist researchers and engineers in making valid assumptions when using forward/reverse paths. In addition, we contribute to give a conclusive overview on the partial asymmetry of Internet routing.

The usability of Traceroute for measuring reverse paths is, depending on the application, questionable. However, for mitigating DDoS attacks closer to the source we have identified atleast one additional opportunity. We have confirmed

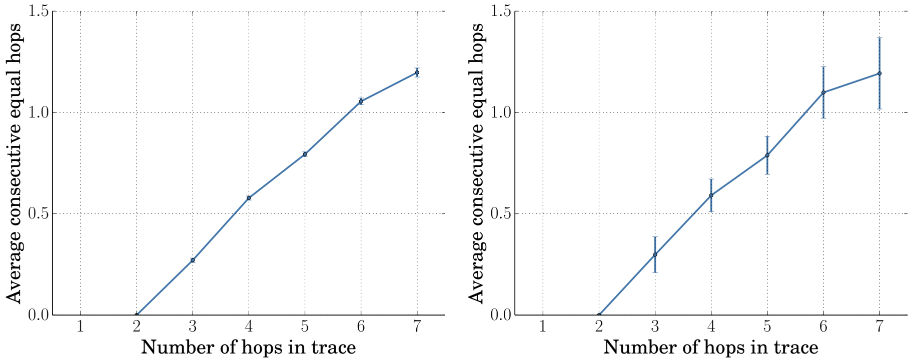

(a) All results

(b) First result per pair only

Fig. 8: Average number of CEH between forward and reverse paths

the presence of asymmetry in the majority of Internet routes, and determined where this asymmetry occurs. Our hypothesis, that DDoS attacks can be mitigated outside the victims network using network paths measured using Traceroute has been confirmed. We have found, in the worst case, a hop, representing an AS, is the same in the forward and the reverse path in 59% of the cases, but often more. In addition, we have found that Tier 1 providers are responsible for a large part of Internet Asymmetry.

As future work we plan to extend the analysis on the IP-level. Furthermore, we plan to apply machine learning to estimate network path accuracy given certain indicators, such as the type of networks that are involved and the length of the path.

Acknowledgments The authors would like to thank RIPE Atlas for facilitating the measurements on their network. RIPE Atlas is a global network of thousands of probes that measure Internet connectivity and reachability, providing an unprecedented understanding of the state of the Internet in real time. RIPE Atlas is coordinated by the RIPE NCC. In addition, the authors would like to thank Roland van Rijswijk for his insights and ideas.

This work was funded by the Network of Excellence project FLAMINGO (ICT-318488), which is supported by the European Commission under its Seventh Framework Programme.

References

1. Damerau, F.J.: A technique for computer detection and correction of spelling errors (1964)
2. Ferguson, P., Senie, D.: Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing. RFC 2827 (Best Current Practice) (2000), `http://www.ietf.org/rfc/rfc2827.txt`

3. He, Y.H.Y., Chen, W.C.W., Xiao, B.X.B., Peng, W.P.W.: An Efficient and Practical Defense Method Against DDoS Attack at the Source-End. 11th International Conference on Parallel and Distributed Systems (ICPADS'05) 2 (2005)
4. He, Y., Faloutsos, M., Krishnamurthy, S., Huffaker, B.: On routing asymmetry in the Internet. In: Global Telecommunications Conference, 2005. GLOBECOM '05. IEEE. vol. 2, pp. 6 pp.– (Nov 2005)
5. He, Y., Faloutsos, M., Krishnamurthy, S.V.: Quantifying routing asymmetry in the Internet at the AS level. In: GLOBECOM. pp. 1474–1479. IEEE (2004), `http://dblp.uni-trier.de/db/conf/globecom/globecom2004.html#HeFK04`
6. Iputils: ping(8) (2014), `http://man7.org/linux/man-pages/man8/ping.8.html`
7. Katz-Bassett, E., Madhyastha, H., Adhikari, V.: Reverse traceroute. In: Network Systems Design and Implementation (2010), `https://www.usenix.org/legacy/event/nsdi10/tech/full_papers/katz-bassett.pdf`
8. Levenshtein, V.I.: Binary codes capable of correcting deletions, insertions, and reversals. Soviet Physics Doklady 10, 707–710 (1966)
9. Mirkovic, J., Prier, G., Reiher, P.: Attacking DDoS at the source. In: Proceedings - International Conference on Network Protocols, ICNP. pp. 312–321 (2008)
10. Paxson, V.: End-to-end Routing Behavior in the Internet. In: Conference Proceedings on Applications, Technologies, Architectures, and Protocols for Computer Communications. pp. 25–38. SIGCOMM '96, ACM, New York, NY, USA (1996), `http://doi.acm.org/10.1145/248156.248160`
11. Prince, M.: Technical Details Behind a 400Gbps NTP Amplification DDoS Attack. `http://blog.cloudflare.com/technical-details-behind-a-400gbps-ntp-amplification-ddos-attack`, accessed on 19 Februari 2014
12. Prince, M.: The DDoS That Almost Broke the Internet. `http://blog.cloudflare.com/the-ddos-that-almost-broke-the-internet`, accessed on 27 January 2014
13. Schwartz, Y., Shavitt, Y., Weinsberg, U.: On the Diversity, Stability and Symmetry of End-to-End Internet Routes. In: INFOCOM IEEE Conference on Computer Communications Workshops , 2010. pp. 1–6 (2010)
14. Simon, D.R., Agarwal, S., Maltz, D.A.: AS-based accountability as a cost-effective DDoS defense. Proceedings of the first conference on First Workshop on Hot Topics in Understanding Botnets p. 9 (2007), `http://portal.acm.org/citation.cfm?id=1323128.1323137`