creative commons

COMMONS DEED

저작자표시-비영리-변경금지 2.0 대한민국

이용자는 아래의 조건을 따르는 경우에 한하여 자유롭게

- 이 저작물을 복제, 배포, 전송, 전시, 공연 및 방송할 수 있습니다.

다음과 같은 조건을 따라야 합니다:

저작자표시. 귀하는 원저작자를 표시하여야 합니다.

비영리. 귀하는 이 저작물을 영리 목적으로 이용할 수 없습니다.

변경금지. 귀하는 이 저작물을 개작, 변형 또는 가공할 수 없습니다.

- 귀하는, 이 저작물의 재이용이나 배포의 경우, 이 저작물에 적용된 이용허락조건을 명확하게 나타내어야 합니다.
- 저작권자로부터 별도의 허가를 받으면 이러한 조건들은 적용되지 않습니다.

저작권법에 따른 이용자의 권리는 위의 내용에 의하여 영향을 받지 않습니다.

이것은 이용허락규약(Legal Code)을 이해하기 쉽게 요약한 것입니다.

Disclaimer

dCollection

Ph.D. DISSERTATION

Internet AS-level Topology: Discovery and Analysis

인터넷 AS-Level 토폴로지: 발견과 분석

AUGUST 2014

DEPARTMENT OF ELECTRICAL ENGINEERING AND COMPUTER SCIENCE
COLLEGE OF ENGINEERING
SEOUL NATIONAL UNIVERSITY

Akmal Khan

Ph.D. DISSERTATION

Internet AS-level Topology: Discovery and Analysis

인터넷 AS-Level 토폴로지: 발견과 분석

AUGUST 2014

DEPARTMENT OF ELECTRICAL ENGINEERING AND COMPUTER SCIENCE
COLLEGE OF ENGINEERING
SEOUL NATIONAL UNIVERSITY

Akmal Khan

Internet AS-level Topology: Discovery and Analysis

인터넷 AS-Level 토폴로지: 발견과 분석

지도교수 권태경

이 논문을 공학박사학위논문으로 제출함

2014 년 5 월

서울대학교 대학원

전기 · 컴퓨터 공학부

칸 아크말

칸 아크말의 박사학위논문을 인준함

2014 년 6 월

위 원 장	김 종 권	(인)
부위원장	권 태 경	(인)
위 원	최 양 희	(인)
위 원	최 성 현	(인)
위 원	백 상 헌	(인)

Abstract

The Autonomous System (AS) level topology of the Internet is critical for future protocol design, performance evaluation, simulation and analysis. Despite significant research efforts over the past decade, the AS-level topology of the Internet is far from complete. Worse, recent studies highlight that the incompleteness problem is much larger than previously believed. In this thesis, we highlight the importance of two under utilized AS-level topology data sources: Looking glass (LG) servers and Internet Routing Registries (IRR).

By querying Looking glass (LG) servers, we build an AS topology estimate of around 143 K AS links from 245 LG servers across 110 countries. We find 20 K new AS links in the AS topology from the LG servers. We observe 620 neighboring ASes of the LG servers that are not sharing their BGP traces with any of RouteViews [49], RIPE-RIS [65], and PCH [66]. We discover 686 new ASes in the AS topology from the LG servers that are hidden from other AS topologies. Overall, we conclude that collecting BGP traces from the LG servers help increase the narrow view of BGP observed from current BGP collectors [38]. However, the AS topology view from the LG servers suffers from limited vantage points of the LG servers and BGP export policies employed by the neighboring ASes of LG servers.

Understanding the benefits and limitations of LG servers, we explore Internet Routing Registries (IRR), which are a set of databases used by ASes to register their inter-domain routing policies. More specifically, we first present a methodology to extract AS-level topology (e.g., bilateral and multilateral peering links) from the IRR. We extract 610 K AS links from the IRR dataset of Nov. 1st, 2013; 68% of which can be matched in BGP, traceroute, and in the cliques of Internet eXchange points (IXPs). We find active usage of the IRR by member ASes of IXPs, which results in inferring peering matrices of many large and small IXPs. Finally, we present a methodology to infer business relationships between ASes using routing polices stored in the IRR. We show that the overall accuracy of our algorithm is comparable (97% for p2c, 95% for p2p links) to the existing algorithms, which infer AS relationships using BGP AS paths. We conclude that the IRR is a strong complementary source for better understandings of the structure, performance, dynamics, and evolution of the Internet since it is actively used by a large number of operational ASes in the Internet.

Keywords: Inter-domain Routing, Looking Glass (LG) Servers, Internet Routing Registry (IRR)

Student Number: 2008-23532

Contents

Abstract **i**

Contents **iii**

List of Figures **v**

List of Tables **vii**

Chapter 1 Introduction **1**

Chapter 2 Background **6**

2.1 Inter-domain Routing 6

2.2 Importance of Research on AS topology 8

2.3 Looking Glass Servers 13

2.4 Internet Routing Registries 15

2.5 Related Work 17

Chapter 3 METHODOLOGY **20**

3.1 AS Topology derived from LG servers 20

3.2 Exploring IRR for AS-level Topology 22

3.2.1 IXPs (IP Prefixes, ASNs, and Members) 22

3.2.2 Route Servers (ASNs and AS-Set Objects) 23

3.2.3 Preprocessing IRR data 23

3.2.4 Extracting AS Links and Policies from IRR 27
3.3 AS Relationship Inference 30

Chapter 4 Datasets **32**
4.1 AS Topologies 32
4.2 AS Relationship Datasets 34

Chapter 5 Analysis **36**
5.1 Comparison of BGP feeders 36
5.2 Registration of Routing Policies in the IRR 39
5.2.1 Policies in aut-num Objects 39
5.2.2 Number of Local Preference (LocalPref) Values in the IRR 42
5.3 Analysis on AS-level Topology 43
5.3.1 Overlapping and Missing IRR-based AS Links 43
5.3.2 BGP-based AS Links vs. IRR-based AS Links 48
5.3.3 AS Degree Distribution 50

Chapter 6 AS Relationship Inference **53**
6.1 Evaluation Results 53

Chapter 7 Summary & Future Work **55**

요약 **65**

Acknowledgements **67**

List of Figures

Figure 2.1 Example of RPSL policy for AS1. 19

Figure 3.1 Year in which aut-num and as-set objects are updated, as in the IRR dataset of Nov. 1st, 2013. 75% of the total routing policies in aut-num objects have been updated since 2012. 25

Figure 3.2 AS-level links and annotations extracted from a sample policy of AS1. The lines with solid arrows represent customer-to-provider (c2p) links and dotted lines represent peer-to-peer (p2p) links between ASes with the annotation showing a local preference (LocalPref) value or a peering location (e.g., DE-CIX). 28

Figure 5.1 Policy registration practice of ASes in IRR. 40

Figure 5.2 Import/export routing policies in aut-num objects are registered differently per AS. 40

Figure 5.3 The distributions of the number of Local preference (LocalPref) values registered by individual ASes in the IRR are shown based on PeeringDB [62]. 41

Figure 5.4 Fraction of per-AS IRR Links missing in BGPAll is shown depending on regions and IRR update year. 45

Figure 5.5 per-AS BGP Links overlapping in IRR. 51
Figure 5.6 AS Cumulative degree distribution (PDF) of the AS topologies extracted from BGP (BGPAll), traceroute (TrAll), IRR and verified one (BGPAll+IRRverified). 52

List of Tables

Table 2.1 A sample result of the `show ip bgp summary` command. . 14

Table 2.2 A sample result of the `BGP neighbor ip advertised routes` command. 15

Table 3.1 An example of AS relationship inference. 30

Table 5.1 The number of overlapping and unique (in bold) ASes and peering routers between various BGP feeders (Route-Views (RV), RIPE-RIS, PCH, and LG servers). 38

Table 5.2 The portion of IRR AS links that are overlapping with other datasets is shown. 43

Table 5.3 IRR AS links overlapping in BGPAll, IXPAll, and CombinedAll datasets. 47

Table 5.4 Top 15 IXPs (in terms of member ASes reported by Isolario) and their peering matrices in IRR and a Combined (BGP, Traceroute, and IRR) dataset. 49

Table 6.1 IRR AS Relationships compared with others. 54

Chapter 1

Introduction

The Internet consists of thousands of Autonomous Systems (ASes) that exchange inter-domain routing information using Border Gateway Protocol (BGP) [1]. The entire Internet can be viewed as an AS-level topology graph where each AS is a node, and a BGP connection between two ASes is a link. The importance of the AS topology has been highlighted through many studies, such as analyzing Internet topological properties [20, 21, 33], inferring AS relationships [14], building network topology generators for simulations [35], and evaluating the effectiveness of new protocols and improvements [13]. Considering the importance of the AS topology in many areas of networking research, significant efforts [12, 21, 24, 26, 29, 39, 40, 61] have been made to discover and construct it. However, it still remains as a challenge to develop a complete and accurate view of the AS-level topology [12, 21, 24, 30, 38]. Worse, recent studies [30, 31] highlighted that the incompleteness problem of AS-level topology is much severe than previously recognized. For example, Giotsas *et al.* [31] use BGP community values to infer 206 K peer-to-peer (p2p) links from 13 large European Internet eXchange points (IXPs), four times more number of p2p links than what can be directly observable in public BGP data. Still, their approach has a limitation since BGP community values are not used by all ASes in the Internet.

One of the promising data source for discovering the AS-level topology is a Looking glass (LG) server, which is a web based portal operated by network operators (e.g., ISPs and NOCs) to provide a look into the BGP routing tables of the ASes in which the server resides. The importance of LG servers in constructing the Internet AS-level topology has been highlighted in many studies [21,24,29]. For instance, Augustin *et al.* [29] uses the `show ip bgp summary`[1] and `traceroute` commands to map IXPs, their members, and their peering matrices. While prior studies have shown the usefulness of LG servers, it is not clear what other information (apart from those available with the `show ip bgp summary` command) is available with LG servers for the purpose of collecting the AS topology. Thus, we conducted a comprehensive investigation to find out how many LG servers are operational and what functionalities are provided by individual LG servers.

Another inter-domain routing data source for exploring the AS-level topology is the Internet routing registries (IRR), which are a set of databases used by ASes to register their inter-domain routing policy information. Earlier research has highlighted the role of IRR in discovering AS-level topology of the Internet [10, 11, 21]. However, none of the publicly available AS topology datasets [24, 60, 75] contain AS links observed in the IRR [27, 36], primarily due to the negative beliefs that IRR often contains incomplete or outdated information [23, 24]. While the quality of routing policies registered in the IRR by many ASes has greatly improved in recent years [5], there has been no study that empirically observes: (1) how many ASes are registering their routing policies in the IRR, (2) how routing policy registration practices of ASes in the IRR varies across different Regional Internet registries (RIR) regions, (3) how the registered routing policies can be used for extracting AS links (e.g., bilateral

[1] `show ip bgp summary` lists the BGP sessions established with an LG router, and details the ASN and IP address of its peering BGP router, for each BGP session.

and multilateral peering links at IXPs) and inferring AS relationships, and (4) how many AS links extracted from the IRR can be observed in BGP and vice versa. Finding empirically-grounded answers to these questions is important as it helps the research and operational community understand whether and how helpful the IRR data are for AS topology construction as well as AS relationship inference.

In this thesis, we propose a comprehensive methodology to discover AS links by querying LG servers and extract AS links from the routing policies of ASes registered in the IRR. More specifically, we design, implement, and evaluate tools to study Looking glass (LG) servers and Internet Routing Registries (IRR): (i) We implement and active measurement tool which given the name of an LG server and a BGP command (e.g., BGP summary command) connects to the LG servers over the Internet and run multiple queries and collects the results in the form of html output from the LG server and stores the results in MySQL database for later parsing the results and extracting AS topology and other link annotations. Note that, in our evaluation, we have only shown the results of running BGP commands on the LG servers. However, this tool can also be used to run ping/traceroute from the thousands of LG servers sites distributed all over Internet for running real time active measurements over the Internet (e.g., troubleshooting purposes, detecting path changes, etc.). (ii) We implement a comprehensive IRR routing policy parser which given a routing policy of an AS in the IRR can extract stored AS topology information and other link and policy annotations. For example, our tool extracts bilateral and multilateral peering links of ASes at IXPs. Second, we compare the AS topology obtained from the IRR against the ones observed in BGP (publicly available BGP traces and BGP traces collected from LG servers), traceroute, and the cliques of IXPs [21]. Finally, we propose a method to infer business relationships (e.g., peer-to-peer or provider-to-customer) between ASes using routing polices registered in the IRR. Our proposed method is complementary to the existing ones (e.g., CAIDA [22], Isoalrio [25], Gao [9], etc.) that rely on BGP AS paths to infer AS relationships. The key insights made with querying LG servers in the month of Oct. 2013 and the IRR datasets of Nov. 1st, 2013, are as follows.

1. We collect around 143 K AS links from 245 LG servers across 110 countries. We find 20 K new AS links in the AS topology from the LG servers. We observe 620 neighboring ASes of the LG servers that are not sharing their BGP traces with any of RouteViews [49], RIPE-RIS [65], and PCH [66]. Overall, we conclude that collecting BGP traces from the LG servers help increase the narrow view of BGP observed from current BGP collectors [38]. However, the AS topology view from the LG servers suffers from limited vantage points of the LG servers and BGP export policies employed by the neighboring ASes of LG servers. (§ 3.1)

2. We find 26,657 ASes (47% of 56,718 allocated ASNs) are registering their routing policies in the IRR. We observe 92% ASes in RIPE and 50% ASes in APNIC region are registering their routing policies in the IRR. On the other hand, only 4-16% of ASes in LACNIC, ARIN and AfriNIC regions register routing policies in the IRR. We also observe 16 K ASes are registering their BGP local preference values (ranging from 1 to 11) in the IRR. We find that various types of ASes (e.g., content, Network Service Provider) are using different levels of local preference values for the purpose of traffic engineering (§ 5.2).

3. We extract an AS-level topology snapshot consisting of 54 K ASes and 610 K links from the IRR. We find 68% of 610 K AS links obtained from the IRR are matched in BGP, traceroute, and the IXP cliques. We also find that member ASes of large and small IXPs are actively registering their policies in the IRR, which helps in discovering 295 K peer-to-peer relationships in the IRR that match with the IXP cliques. (§ 5.3)

4. We propose a new IRR-based AS relationship inference method and evaluate it against the four existing algorithms (e.g., Gao [9], Cyclops [24], Isolario [25], and CAIDA/Luckie [22]) and two ground truth datasets provided by Luckie *et al.* [22]. We find that the overall accuracy of our AS relationship inference method is comparable (97% for provider-to-customer and 95% for peer-to-peer) to the existing approaches, including the most recently proposed one by Luckie *et al.* [22] (§ 6).

5. Based on our findings, we advocate the role of LG servers and the IRR in better understanding the structure, performance, dynamics, and evolution of the Internet since LG servers provides a near real time view into an operational network and IRR is actively used by a large number of operational ASes in the Internet.

We organize this dissertation as follows. In Chapter 2, we present the background on inter-domain routing, Looking glass servers, and the IRR. Chapter 3 describes our methodology on (i) how to query LG servers , (ii) how to extract AS links information from the IRR, and (iii) how to infer AS relationships from routing policies registered in the IRR. Chapters 5 and 6 compare the IRR-based AS-level topology and the existing other topologies (e.g., BGP-based or Traceroute-based) and evaluate our method proposed to infer AS relationships, respectively. We conclude the dissertation in Chapter 7.

Chapter 2

Background

In this chapter, we briefly describe an overview of Internet routing, Looking glass servers, Internet Routing Registries and the Routing Policy Specification Language.

2.1 Inter-domain Routing

Autonomous Systems (ASes) use Border Gateway Protocol (BGP) to define their routing policy that help control which IP prefixes or routes are chosen and which routes are to be propagated to their neighbor ASes [8]. Each route is tagged with a number of BGP attributes such as **AS-PATH**, which is a sequence of ASes between source and destination ASes. The **LocalPref** (local preference) attribute indicates the degree of preference of one route over the other routes. The **Community** attribute is a 32-bit integer used to influence the routing policies of the provider AS such as for traffic engineering purposes. **Multi-Exit Discriminator (MED)** is assigned to a route to determine the exit point to a destination AS. Each AS is identified by a number (ASN), which is assigned by Regional Internet Registry (RIR)[1].

Business Relationships: The connectivity between ASes do not imply reachability, as whether the traffic from a certain AS in the Internet can reach another AS depends on its business relationships with its neighboring ASes. In general, an AS can have three types of relationship to its neighbor ASes [10]. The **customer-to-provider (c2p)** type (or **provider-to-customer (p2c)**, if looked at from the opposite direction) in which the customer AS buys transit access to the Internet from a provider AS. A **peer-to-peer (p2p)** relationship is established when two ASes agree to exchange traffic between each others' ASes, typically for free. Finally, **sibling-to-sibling (s2s)** type relationships are used between ASes operated by the same organization, where each AS may serve a different role (e.g., backbone, regional networks, etc.). More complex relationships can exist, e.g., backup links, and variations of the simple relationships that are described above such as partial transit and paid peering [4].

Peering Agreements and Policies: There are generally two ways in which ASes connect (also known as BGP peering) with other ASes [4]: (i) When two networks negotiate with each other and establish a peering session directly, we call it **Bilateral peering**. ASes can setup a bilateral peering session at an IXP, which is then called **Public peering**. Alternatively, **Private peering** (Private Network Interconnect (PNI)) is a direct interconnection between two ASes for exchanging a large volume of traffic, using a dedicated transport service or fiber. (ii) While bilateral peering offers the most control, some networks with very open peering policies may wish to simplify the process, and simply "connect with everyone". To help facilitate this, many IXPs offer **Multilateral peering agreements (MLPA)**[2].

Peering policies of an AS suggests which ASes it can peer with or not.

[1]The five operational RIRs are AfriNIC [70], APNIC [71], ARIN [72], LACNIC [73], and RIPE NCC [74].

[2]IXPs manage MLPAs using route servers (RS) which allow member ASes to establish a single BGP session with a route server and receive routes from every other AS connected to the route server.

According to its business requirements, an AS can have four types of peering policies: **Open Peering** policy implies that an AS is willing to peer with anyone (except its customers) without imposing specific conditions, while **Selective Peering** policy means that an AS is generally willing to peer with those who meet a specific set of conditions. **Restrictive Peering** policy means that an AS is generally inclined not to add any new peers; typically used by Tier-1 transit providers, and **No Peering** policy means that an AS does not peer at all, as it is interested in selling transit services.

AS Topology Data Sources: Since their is no single authority running the Internet, no single data source has a complete topology of the Internet. There have been three main approaches to construct the AS topology, each of which has its own limitations: (i) Passive measurements by collecting BGP routing tables and updates suffer from routing policy filters and best path selection decisions made by neighboring ASes of BGP collectors [26, 38]. (ii) Active measurements using traceroute are error-prone and generate potentially false AS links due to non-responsive hosts and errors in converting IP addresses to AS numbers (IP-to-AS mapping) [26, 37]. (iii) Internet Routing Registries (IRRs) are believed to contain outdated information, as AS links extracted from the IRRs can be outdated or not yet operational [21]. However, this thesis shows that the negative beliefs regarding the IRRs do not hold as much as they possibly did a decade ago [5]. Consequently, we show that AS topology observed through BGP can be significantly augmented using the information extracted from the Internet Routing Registries.

2.2 Importance of Research on AS topology

The importance of research on AS topology has been highlighted through many studies [13,14,20,21,21,33,35,78,83–88]. We give a brief overview of four of the important research areas which require the AS topology for better understanding and/or improving some aspects of the Internet: (1) Analysis of Internet topological properties, (2) Evaluation of new routing protocols, (3) Solving security issues of Inter-domain routing, and (4) AS relationship inference.

1. Analysis of Internet Topological Properties: The AS topology snapshots are often used to study the graph theoretic properties of Internet. For example, He *et al.* [21] studies the effects of peer-to-peer edges on Internet topological properties such as path lengths. They report that for some ASes more than 50% of the paths stop at provider ISPs assuming policy-aware routing. They also report dramatic changes observed in some of the results reported in studies that have used incomplete AS topology snapshots, e.g., routing decisions and ISP profit/cost. Govindan *et al.* [78] studies routing stability, i.e., transient changes in routes caused by router and link failures or router misconfigurations. Such analysis helps in a better understanding of how route stability impact end-to-end communication performance. Link-Rank [87] annotates links in AS topology with weights, which are calculated based on number of routes using that link. Such link weight annotations helps in detecting various kinds of routing problems. For example, changes in some routes can affect the delay, loss, packet re-ordering, and throughput characteristics observed by long-lived connections.

Gupta *et al.* [88] studies the causes of circuitous Internet paths (also known as path detour/tromboning) and evaluate the benefits of increased peering and better cache proxy placement for reducing latency to popular Internet sites (e.g., Google) in Africa. They observe high network latencies to popular destinations due to circuitous Internet paths, i.e., paths that should remain local in Africa are being observed to detour through Europe. The main reason for such high latencies is due to connectivity of ASes in Africa, i.e., due to non-existent peerings between some large ASes in Africa, even the increased placement of Google caches do not result in decreasing network latencies. Path detouring issues are also observed in many other regions (e.g., US) due to Internet peering issues between ASes. Thus, a complete view of the AS topology can help guide

ISPs in making better peering decisions which in turn can help in reducing the network latencies.

2. Evaluation of new routing protocols: The Internet is the largest and most popular man made network infrastructure, which is used by billions of people and organizations on daily basis to execute their social and business related tasks. To communicate over the Internet, applications (e.g., web surfing) only needs to know the IP address of the communicating entity, which can be retrieved using DNS (i.e., website URL to IP mapping). As the usage of the Internet has been continuously increasing, since its commercialization in the 90s, so are the concern over the various security related issues which were absent at the time of design of Internet architecture.

While there are various network security related issues, how packets are routed between source and destination entities becomes an increasingly important issue. For instance, a source may want to block her packets from going through internet service providers (ISPs) that are suspicious of wiretapping or censorship (e.g., by government). We have little visibility (not to mention, control) over packet routing. That is, source entities have little idea about which autonomous systems (ASes) are participating in forwarding/routing their packets. The motivation of path control are similar for the destination entity, in addition, a destination entity can also be interested in path control due to traffic engineering purposes.

There are numerous proposals debating on the need of source and/or destination controlled routing [83–86]. Under NIRA [85], for instance, senders choose the path into the Internet core, and receivers choose the path out. Similarly, Pathlets [84] allows senders to choose paths and providers specify policies based on the previous hop in the path. SCION [83] proposes secure AS-level route control by using a hierarchical trust relationship among ASes. ICING [86] suggests a strong path verification mechanism by assuming the existence of a consent server for each node (e.g., a router or an AS) and the setup of a shared key for every pair of nodes. In ICING, a source should find out (and select) which nodes to visit toward a destination and contact the contact servers of the selected nodes, which may not be feasible in the near future. Overall, an AS topology view is required to enable applications on hosts to select end-to-end path selection and the performance evaluation of the above mentioned protocols also require the AS topology.

3. Solving Security Issues of Inter-domain Routing: BGP is vulnerable to misconfigured and malicious routing information as there is no verification mechanism of the incoming routing information. One of the most notorious BGP attack is IP prefix hijacking, which occurs when a malicious or misconfigured BGP router originates an IP prefix that the router (or the AS that contains the IP subnet) does not own. IP prefix hijacking is essentially a special form of denial of service attack. Even though BGP operates well in practice due to simplicity and resilience, some outages may lead to significant and widespread damage. For instance, one of the early BGP hijacks happened in 1997, where traffic to be redirected to AS7007 hijacked a lot of specific (or longer) IP prefixes. Some of the more recent incidents of that kind are ConEd (in 2005) and an outage for the popular YouTube site caused by Pakistan Telecom in 2008. As the number of critical applications (online banking, stock trading, and telemedicine) on the Internet grows, there will be more dependency on the underlying network infrastructure to provide reliable and secure internet connectivity.

Research community on inter-domain routing has worked out many protocols and technical contributions for BGP operational issues such as scalability, convergence, routing stability, and performance. However, the security aspects of BGP have not been practically solved. There has been a large body of research on routing security [23, 81] to ensure the authenticity and correctness of topology propagation and route computation. For instance, BGPSEC [89] improves security for BGP routing. BGPSEC relies on Resource Public Key Infrastructure (RPKI) [41] whose deployment has already been started. However, it will take many more years before a full scale deployment of BGPSEC is expected [81] and attacks are still possible even after the deployment of RPKI/BGPSEC [90]. Thus, to detect various types of prefix hijacking and traffic interception attacks, several passive and active measurement systems have been proposed. The AS topology is an important component of most of the proposed systems as it helps to detect invalid paths [23,81]. For example, Krugel *et al.* [91] proposed to gather route validation information through BGP traces to identify prefix hijacking attacks. They proposed to use the AS topology view to check an AS path validity, e.g., in a valid path, two neighboring ASes should be in the same geographic region and the path should traverse the core at most once. Certainly such an approach requires a complete AS topology view to detect AS path validity since an incomplete AS topology view can miss many operational paths in the Internet.

4. AS Relationship Inference: Accurate knowledge of business relationships between ASes is relevant to both technical aspects (e.g., network robustness, traffic engineering) and economy-based modeling of the evolution of Internet [22]. However, as business relationships between ASes are generally not publicly disclosed, considerable effort has been made to infer AS relationships between ASes. The most complete AS topology is an important part of inferring accurate AS relationship inference [22] as inability to observe some AS paths (e.g., peer-to-peer links) can result in inferring inaccurate AS relationships.

The seminal work by Gao [9] infers relationships between ASes based on the valley-free property of AS paths, i.e., each AS path consists of an uphill segment of zero or more c2p or sibling links, zero or one p2p links at the top of the AS path, followed by a downhill segment of zero or more p2c or sibling links. More recently, Luckie *et al.* [22] proposed a method based on less restrictive valley-free property rules and validated 34.6% of their inferred AS relationships. Most of the proposed AS relationship inference methods [9, 22] use AS paths observed in BGP. In contrast, similar to Nemecis [10], we highlight that the information stored in the Internet Routing Registries also can be used to infer AS relationships accurately.

2.3 Looking Glass Servers

Looking glass (LG) servers are web based portals operated by network operators to provide a look into the BGP routing tables of the ASes in which the server resides. For example, from a response of a query to an LG server, a network problem can be traced back to its reasons like misconfigured BGP route advertisement, wrong route aggregation, or misconfigured AS path prepending. Traditionally, such accesses to route collectors have been provided through Telnet. However, many networks are currently operating LG servers instead of providing Telnet access to their BGP routers directly [3]. By an LG server, we mean a web site that allows running commands (e.g., traceroute) from one or more BGP routers that are under the control of the LG server. For instance, the LG server of Hurricane Electric[4] provides facilities to run LG commands on its BGP routers that are distributed across 92 locations worldwide.

The importance of LG servers in constructing the Internet AS-level topology has been highlighted in many studies [21, 24, 29]. For instance, Augustin *et al.* [29] uses the `show ip bgp summary`[5] and `traceroute` commands to map IXPs, their members, and their peering matrices. While prior studies have shown the usefulness of LG servers, it is not clear what other information (apart from those available with the `show ip bgp summary` command) is available with LG servers for the purpose of collecting the AS topology. Thus, we conducted a comprehensive investigation to find out how many LG servers are operational and what functionalities are provided by individual LG servers. We first build a list of LG servers from the following sources: peeringDB [62], Tracer-

[3] A BGP router under the control of an LG server is called an LG router.

[4] Hurricane Electric LG. http://lg.he.net

[5] `show ip bgp summary` lists the BGP sessions established with an LG router, and details the ASN and IP address of its peering BGP router, for each BGP session.

Table 2.1 A sample result of the `show ip bgp summary` command.

Router: cr1-eqx3-pa3 Local AS Number: 29075					
Command: show ip bgp summary					
Neighbor	**AS#**	**State**	**Time**	**Received**	**Sent**
195.42.144.104	6939	ESTAB	61d	36,464	153

oute.org [42], Traceroute.net.ru [43], BGP4.as [44], BGP4.net [45], and Virus-net [46]. After removing the overlapping LG servers from the above sources, we find 1.2 K LG servers, only 420 of which were in operation at the time of this study, in the month of March 2013. Our scripts can query 388 LG servers since the web sites of the other 20 LG servers are not parsable and 12 LG servers limit automated queries.

We queried 388 LG servers (running on 410 ASes) to learn their supported functionalities. We find that as many as two dozen commands are supported by different LG servers, while a few of them are more widely supported than others. For example, all the 388 LG servers support `traceroute` and `ping` commands from 4.4 K (in total) locations in the Internet. Another widely supported command is `show ip bgp summary`, which is supported by 245 LG servers from 1.9 K locations. The regional Internet registries (RIR) wise distribution of 245 LG servers are as follows: RIPE (175), ARIN (40), APNIC (15), LACNIC (13), and AfriNIC (2).

Table 2.1 illustrates a sample result of querying a router (cr1-eqx3-pa3 operating at Paris Equinix) with the `show ip bgp summary` command through the LG server provided by Ielo (AS29075). It shows that Ielo has a BGP session with Hurricane Electric (AS6939) at Paris Equinix. It also shows other important information, such as (i) how long the BGP session has been alive (61 days), (ii) 36,464 routes received from the BGP neighbor, and (iii) 153 routes advertised to the BGP neighbor over this link.

We also find that 59 LG servers (distributed over 250 locations) allow us to run the `BGP neighbor ip advertised routes` command, which helps observe IP prefix announcement(s) advertised by an LG router to its peering BGP

Table 2.2 A sample result of the `BGP neighbor ip advertised routes` command.

Router: cr1-eqx3-pa3 Local AS Number: 29075		
Command: BGP neighbor 195.42.144.104 advertised routes		
Prefix	**Next Hop**	**AS PATH**
149.154.80.0/21	195.42.144.71	29075 50618 57141
91.227.48.0/24	195.42.144.71	29075 50618 25091 56728

routers. Table 2.2 shows a sample result of the `BGP neighbor ip advertised routes` command on the BGP router cr1-eqx3-pa3. Each row shows an IP prefix, its next hop address and AS path information.

2.4 Internet Routing Registries

The Internet Routing Registries (IRRs) are a set of databases storing routing policy information of ASes, such as IP prefixes originated by ASes and routing policies towards their neighbor ASes [6]. There are numerous IRRs maintained by large ISPs (Level3, NTT), small ISPs (Verio), non-affiliated (RADb, AltDB), and Regional Internet Registries (RIRs). While a large number of IRRs (33 in total as of Nov. 1st, 2013) are mirrored at the IRR site [2], RIPE [56], AfriNIC [58], and APNIC [59] IRRs are only available from their own FTP servers. The IRRs of Korea Network Information Center (KRNIC), Japan Network Information Center (JPNIC), Taiwan Network Information Center (TWNIC), and Indian Registry for Internet Names and Numbers (IRINN) are also mirrored at the APNIC FTP server [59]. In total, we collected publicly shared daily snapshots of the whole 40 IRR datasets in the period of Oct 1st, 2010 to Nov. 1st, 2013. Hereafter, we call the whole combined dataset as the IRR dataset.

Routing Policy Registration in the IRR: When registering routing policy information in the IRR, a standard language called Routing Policy Specification Language (RPSL) [3] is used. The RPSL defines several kinds of objects,

[6] Internet Routing Registries are also commonly referred to as WHOIS databases.

most of which can be classified into the following three groups: (i) *inetnum or inet6num* objects describe IPv4 or IPv6 address allocation, (ii) *route, route6, aut-num, route-set, as-set* objects describe routing policies, and (iii) *mntner, person, and role* objects describe who administer the routing policies and so on. We briefly describe the details of RPSL objects, which are used to infer and characterize routing policies of ASes in the IRR: A **mntner** object is used to register an authorized entity to add, delete, or modify objects related to an AS. Once a mntner object is created, the maintainer can register RPSL objects of other types; When registering IP prefixes or routes of an AS, **route** objects are used. When an AS needs to create and specify routing policies for a set of neighboring ASes, **as-set** and **route-set** objects are used. The as-set and route-set objects are hierarchial in nature, as they can refer to other as-set or route-set objects, respectively. For registering import and export policies towards neighboring ASes, **aut-num** objects are used.

As RPSL is very flexible, there are many ways to register routing policies in aut-num objects. First an AS can directly use routes like, "from AS9488 import {147.46.0.0/16}". A more convenient way is to group routes using ASN, e.g., "from AS3 action pref=100; accept AS3" means that accept all routes registered by AS3 and assign a local preference (pref [7]) value of 100. Using as-set or route-set is another way, e.g., "from AS2 accept AS2:AS-Customers", which means accept all routes registered by customers of AS2. Less restrictive filters can be used by keywords like **ANY**, which means any routes received. More restrictive filters can be created by combining regular expressions with the aforementioned filters. More details on RPSL can be found in [3].

Figure 2.1 illustrates examples of the aforementioned RPSL objects. RPSL defines several additional attributes that are not shown in the figure, such as the **source** attribute specifying in which IRR the object is registered, and the

[7]Preference (pref) is opposite to BGP local preference (LocalPref) in that the smaller values are preferred over larger values.

changed attribute containing the last updated or created date of an RPSL object.

2.5 Related Work

There have been a number of measurement studies related to the AS topology discovery [12, 21, 24, 29, 39, 60, 61]. To quote the most recent efforts, He *et al.* [21] provide a large scale comprehensive synthesis of the available routing data sources such as BGP routing tables, IRR, and traceroute data. Augustin *et al.* [29] build on the work of He *et al.* [21], but the focus is on the IXP substrate, not on the AS topology as a whole. Active measurement platforms such as Ark [61], DIMES [39], and iPlane [40] are providing the AS topology views, but suffer from the small number of vantage points to run traceroute measurements. To overcome the limitation, Chen *et al.* [26] propose to send traceroute probes from a large number of (992,000 P2P user IPs in 3,700 ASes) P2P clients.

AS topology from LG servers: So far, LG servers have been considered as a secondary source of inter-domain routing data for discovering links in the Internet topology [21,24,29,60]. That is, LG servers have been used to augment some AS links to the AS topology extracted from BGP traces [24], or used to help verify the AS links found in the IRR [21]. To the best of our knowledge, this thesis is the first to show that LG servers are yet another non-negligible source for building Internet AS topology. Moreover, collecting BGP traces from the LG servers can help widen the narrow view of BGP observed from the current BGP collector projects, such as RouteViews, RIPE-RIS, and PCH [38].

AS topology from IRR: IRR has been used to build an AS-level topology of the Internet [10, 11, 29]. However the proposed methods do not consider multilateral peering (MLP) links so that they miss a large number of p2p links in the IRR. Our work finds 389 K p2p links from the IRR, and 295 K of them can be verified in the IXP cliques of hundreds of operational IXPs in the Internet.

Recently, Giotsas *et al.* [31] proposed a method to infer MLP links by using BGP community values in BGP traces as well as by querying LG servers. They inferred 206 K links from 13 large European IXPs. Since BGP community attributes are used by a small number of ASes [31], and information of LG servers are not provided by IXPs, we highlight that IRR can be a complementary source of discovering MLP links as well as other types (e.g., backup links) of missing links, which are not observed in publicly available BGP traces. This thesis further emphasizes the registration practices of stand-alone as-set objects, which are not referenced in the IRR aut-num objects. We show that large number of IRR AS links (e.g., 86 K in the IRR dataset of Nov. 1st, 2013) can be missed if we ignore stand-alone as-set objects.

AS Relationships from IRR: Similar to Nemecis [10], we highlight that the policies stored in the IRR can be used to infer AS relationships. However, unlike Nemecis, we do not rely only on the availability of routing polices of both side ASes of an AS link, to infer their relationship. A similar approach (i.e., relying on the availability of routing policies from both side ASes of an AS link) was used by Luckie *et al.* [22] for possibly more accurate AS relationship inference, which resulted in extracting only 6.5 K p2c relationships from the IRR. In contrast, we show that a larger number of AS relationships for both p2c and p2p types can be accurately inferred even when only one side AS of an AS link has made their routing policies available in the IRR. Since most of other proposed AS relationship inference methods [9, 22] use information of BGP AS paths, we demonstrate that inferring AS relationships from the IRR can help to cross-validate the inferences that are made by BGP AS paths.

```
mntner : Mnt-AS1           mnt-by :   Mnt-AS1
------------------------------------------
route  : 10.1.1.0/16       origin :   AS1
------------------------------------------
as-set : AS1:AS-Customers
members: AS1, AS2:AS-Customers, AS3
as-set : AS2:AS-Customers  members:  AS2, AS10, AS20
as-set : AS1:AS-PEERS-NLIX members:  AS30, AS40
as-set : AS50:AS-RS-Peers  members:  AS1, AS2, AS7, AS8
------------------------------------------
aut-num: AS1
remarks: Customers
import : from AS2 action pref = 50;
           accept AS2:AS-Customers
import : from AS3 action pref = 50;
           accept AS3
export : to AS1:AS-Customers
           announce ANY
remarks: Providers
import : from AS4 action pref = 100;
           accept ANY
export : to AS4
           announce AS1:AS-Customers
remarks: Peer at DE-CIX
import : from AS6 80.81.194.100
           accept AS6
export:  to AS6 80.81.194.100
           announce AS1:AS-Customers
remarks: Peers at NL-IX
import : from AS1:AS-PEERS-NLIX action pref = 80;
           accept ANY
export : to AS1:AS-PEERS-NLIX
           announce AS1:AS-Customers
remarks: Peers at route server
import : from AS50 action pref = 70;
           accept ANY AND NOT AS2
export : to AS50 action community .= {50:50, 50:2};
           announce AS1
```

Figure 2.1 Example of RPSL policy for AS1.

Chapter 3

METHODOLOGY

In this chapter, we first describe our methodology to discover the AS topology from LG servers. Second, we describe how to construct a list of IXPs and route servers, which is needed to extract bilateral and multilateral peering links from the IRR. Finally, we explain methods to extract AS links and infer AS relationships from the IRR.

3.1 AS Topology derived from LG servers

We design a tool to automate a querying process to the 388 LG servers. Our tool issues 30 queries in parallel to the LG servers and waits for 15s between successive queries to the same LG server to avoid overloading them. Collecting data from an LG server is a multi-step process. First, for each LG server our tool learns, by parsing LG server websites, the supported LG commands and its LG routers to which our tool sends queries to collect the data. Second, to each LG router, our tool sends the `show ip bgp summary` command to the LG server. Third, from the returned response of `show ip bgp summary`, our tool extracts IP address(es) of the neighboring router(s) of the LG router. Fourth, by using the IP addresses of the neighboring routers, our tool sends a query of

`BGP neighbor ip advertised routes` to collect the BGP routes advertised by the LG router to its neighboring routers. Finally, all the responses of the `show ip bgp summary` and `BGP neighbor ip advertised routes` commands from the LG server are stored in text files for constructing the AS topology.

We queried 245 LG servers that provide the option of running `show ip bgp summary` command from around 1.9 K locations (distributed across 110 countries), twice a week in the month of Oct. 2013. Total 8 snapshots are combined to create an AS link dataset, which consists of around 70 K AS links. We find 77% of the AS links are intra-AS links, i.e., the source and destination ASes of a link are the same. As we are only interested in inter-AS links in this study, we filter out these intra-AS links and selected only 16 K inter-AS links. Throughout this thesis, AS links refer to those inter-AS links.

We also queried 59 LG servers (out of the 245 ones) that provide the option of running `BGP neighbor ip advertised routes` command, once a week in the month of Oct. 2013. Their LG routers are located in 250 locations distributed across 40 countries. Moreover, these LG servers advertise routes to 5 K routers of their neighboring ASes. From the BGP traces collected from the 59 LG servers, we extracted around 2 million AS paths and broke down these AS paths into around 130 K AS links.

Overall, by running the `show ip bgp summary` and `BGP neighbor ip advertised routes` commands on the LG servers, we have collected 130 K unique AS links (130 K+16 K=143 K-3 K overlapping AS links). To the best of our knowledge, this is the first study that investigates not only `show ip bgp summary` but also `BGP neighbor ip advertised routes` commands to construct the AS topology.

3.2 Exploring IRR for AS-level Topology

3.2.1 IXPs (IP Prefixes, ASNs, and Members)

Since a large number of IXP peering links have been reported in [30,31], we need to check the existence of these peering links in the IRR. To do that, we need a database of IXPs, which includes IXP names, prefixes, ASNs, and their member ASes. However, as is typical for distributed and decentralized systems such as the AS-level ecosystem, there does not exist a publicly available complete and up-to-date centralized database of IXPs. Therefore, we make extensive use of the following four different IXP-related data sources, to collect and synthesize the IXPs' information used in this thesis: (i) In 2009, Augustin *et al.* [29] reported the existence of 359 IXPs, 278 out of which were with a total of 393 known IPv4 prefixes [29]. IXP ASNs and IPv6 prefixes are not contained in the dataset. (ii) Isolario [75] regularly queries web sites of IXPs and publishes the list of 285 IXPs, 221 of which are with member AS information. The IXPs' prefixes and ASNs are not provided here. (iii) PeeringDB [62] contains the information of 450 IXPs along with their IPv4 and IPv6 prefixes. ASN of only one IXP (out of the 450 IXPs) is reported in the dataset. (iv) We searched for the texts like "Internet Exchange" in the description of the IRR aut-num objects, from which we extracted 176 IXP ASNs.

To create a combined list (i.e., removing duplicates) of IXPs from the above four sources, we had to cross-compare IXP's full names and abbreviations. For example, *Deutscher Commercial Internet Exchange* is popularly known as *DE-CIX*. For IXPs with only their IPv4 prefixes available but missing ASN information, we look for the corresponding IXP ASNs that have been registered to be originated by these IXP prefixes in the IRR route objects, or vice versa. Overall, from the four IXP-related data sets, we come up with a list of 570 IXPs along with their prefixes, ASNs, and member ASes (if available).

We rely on the IXP member AS information provided by Isolario [75] wherever possible, since it regularly collects the member information from the respective IXP websites. Yet, we find that the Isolario dataset does not contain member information of many IXPs; even for some very large IXPs such as Equinix, one of the largest operational IXPs in the Internet. In such cases, we next rely on the list of IXP members from PeeringDB.

3.2.2 Route Servers (ASNs and AS-Set Objects)

To extract multilateral peering links (MLP) from the IRR, we first need a list of ASNs being operated as route servers (RS) in the Internet. Since there does not exist a publicly available, complete list of operational route servers, we make a list of 50 route servers ASNs by querying IXP websites, PeeringDB, and the IRR. Moreover, we find as-set objects of these route servers to check their participant ASes. We find as-set objects of 40 IXP route servers by looking for texts like "route server", "-RS", or "ATM"[1] in the description of as-set objects. Note that an IXP can operate multiple ASNs for its operational needs, e.g., AS6777 is the route server ASN of Amsterdam Internet Exchange (AMS-IX) and AS1200 is used for providing other services to its member ASes. Thus, we process routing polices registered (in the IRR) for route server ASNs as MLP ones, while those for IXP ASNs as bilateral peering polices, as will be explained later in section 3.2.4.

3.2.3 Preprocessing IRR data

To extract routing policies from the IRR, we need to preprocess IRR datasets such as for removing duplicate information.

Tagging IRR objects with last-updated date info.

We tag the information in the IRR with the last-updated date to analyze (in section 5.3) whether outdated information in the IRR is the possible reason

[1] Acordo de Tràfego Multilateral (ATM) is the Portuguese acronym for multilateral peering (MLP).

for its mismatch in BGP. To discard possibly outdated information from the IRR, earlier studies [22, 24] have relied on the **changed** attribute of RPSL objects, which contains the last updated date of an RPSL object. We find that we can apply this method on our collected IRR datasets, except RIPE one. Because RIPE, staring from Jan. 2013, has replaced the last updated date of RPSL objects with a dummy date of Jan 1st, 2000, due to data privacy laws in Europe [77]. Yet, we can check whether an object has been updated on a more recent date (than a dummy date of Jan. 1st, 2000) by looking into the historical IRR datasets of ours, which we have collected since Oct. 1st, 2010. For example, we find that the aut-num object of AS29076 (AS-IELO) has been updated on Oct. 15, 2013, i.e., the aut-num object has been far more recently updated than suggested by a dummy date of Jan. 1st, 2000 in the RIPE dataset. Figure 3.1 shows the fraction of the aut-num, as-set, and routing policy entries that have been updated in the year given on the x-axis. We find that 75% of the routing policy entries have been updated since 2012, while the remaining 25% have been updated beforehand. Note that we do not discard these 25% possibly outdated information, as done in other studies [24], since we can not quantify how much of this possibly outdated information, is in actual outdated.

aut-num and as-set objects

As there are multiple Internet Routing Registries (IRRs) in the Internet and different provider or peering ASes may use different IRRs, ASes sometimes are required to register their aut-num objects in multiple IRRs. For example, AS17685 (PLAYONLINE) is registered in 5 registries (NTTCOM, JPIRR, APNIC, JPNIC, and RADb). If an aut-num object is registered in multiple IRRs, we choose to use the most recently updated one, discarding the others.

Filtering aut-num Objects: We find 47,439 aut-num objects in the IRR as of Nov. 1st, 2013; 37,423 objects are registered in one registry, while 4,496 ones are registered in multiple registries. We discard 1,456 aut-num objects

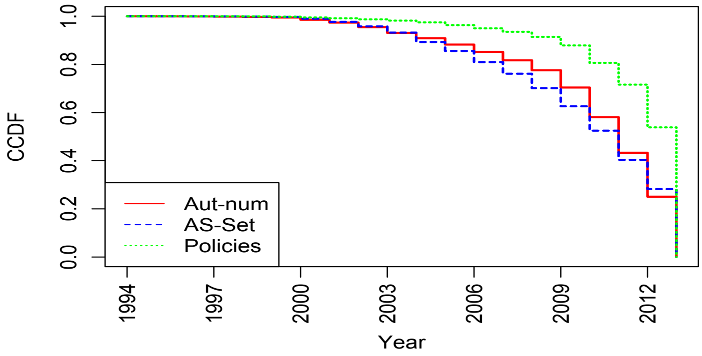

Figure 3.1 Year in which aut-num and as-set objects are updated, as in the IRR dataset of Nov. 1st, 2013. 75% of the total routing policies in aut-num objects have been updated since 2012.

that contain routing policies of ASNs unallocated by the Regional Internet Registries (RIRs). We filter out such aut-num objects. Consequently, we use 40,463 aut-num objects for our study.

Filtering as-set Objects: There are 18,888 as-set objects in our IRR dataset. 18,206 objects are with member ASes, while the rest of them are left empty without such information, thus discarded. We further remove duplicate as-set objects, as done for duplicate aut-num ones. Overall, we use 17,324 as-set objects for our analysis.

as-set to Referrer-AS Mapping

Throughout this thesis, we call an AS referred in an import policy along with an as-set object (as is AS2 with the as-set object AS2:AS-Customers in Figure 3.2, line 12) as a **referrer-AS** of that as-set object. When an as-set object is not appearing in any import policy, then we check whether it has been referred in an export policy, in which case the AS registering the export policy is a referrer-AS (as is AS1 exporting the as-set AS1:AS-Customers in Figure 1, line 17, 20 and 23). This mapping between an as-set object and referrer-ASes is important, since we can infer AS links between an AS exporting an as-set object (i.e., a referrer-AS) and member ASes of the exported as-set objects.

To create an as-set to referrer-AS(es) mapping, we process routing policies of aut-num objects in the IRR dataset of Nov. 1st, 2013. We find referrer-AS(es) of 12,647 (73% of 17,324) as-set objects; 92% of which are referred by only one AS. We observe 1,012 (8%) as-set objects that are referred by more than one referrer-AS, mainly due to the following three (among possibly many) reasons: (a) an as-set object of an IXP can be referred by its members, to find possible peering partners at the IXP. For example, *A*S-UAIX is the as-set object of Ukrainian Internet Exchange (UA-IX), being referred by 68 of its members in their respective aut-num objects. (b) some organizations operate multiple ASes so an as-set object of one AS of the same organization can be referred by other ASes of the organization. For example, *AS-claranet* is referred by 7 ASes (AS8426, AS20869, AS8975, AS8196, AS15385, AS8483, AS6067), all of which are operated by *Claranet*[2]. (c) typographical errors, e.g., *AS-DIGIWEB* is referred by AS31122 and AS3122. On further inspection, we find that AS3122 is a typo and is not a legitimate referrer-AS of AS-DIGIWEB. More specifically, we automate the process of checking and removing such typo by analyzing the frequency with which a referrer-AS is referenced in the registered policies. For example, we find that AS31122 is reported to be a referrer-AS of AS-DIGIWEB in the routing policies of 200 aut-num objects. However, AS3122 is reported to be a referrer-AS in the routing policy of only one aut-num object. Thus, in this case, AS3122 must be a typo of AS31122. We removed 50 instances of such typographical errors from our as-set to referrer-AS(es) mapping.

stand-alone as-set objects: There are no referrer-ASes for the remaining 4,677 (out of 17,324) as-set objects, which are not referred in any of the registered aut-num objects. This can happen as some ASes do not register their policies in aut-num objects but maintain their as-set objects for the operational practices of their neighboring ASes. Thus, for a stand-alone as-set object, we

[2] Claranet. http://noc.eu.clara.net

find an AS who is possibly exporting the as-set object to other ASes, which we call the **exporter-AS** of the as-set object. We find exporter-ASes from the names of as-set objects, as the names often contain their exporter-AS information. For example, the as-set named "AS1:AS-Customers" says it is exported by AS1. However, there are cases where an as-set name contains two or more exporter-ASes, e.g, as in the as-set object "AS1887:AS-Customers:AS12464". In such cases, the rightmost AS is the exporter-AS, following the common practice of as-set registration.

For as-set objects not containing any exporter-ASes in their name, we check whether the maintainer of the as-set object has registered an aut-num object. For example, the maintainer of "AS-PAT-TORIX" is "MAINT-PAT". We find that MAINT-PAT is also maintaining the aut-num object of AS11342. Thus, the exporter-AS of AS-PAT-TORIX is AS11342, which specifies the peerings of AS11342 at Toronto Internet Exchange (TorIX). Overall, we could not find matching exporter-ASes for around 1% of stand-alone as-set objects, as these as-set objects can be mapped to multiple exporter-AS(es). For example, the maintainer of "AS-PTTMETRO-ATM4-SP" operates many ASes for managing IXPs in the LACNIC region. Thus, it is not clear which AS is the exporter-AS of the as-set object "AS-PTTMETRO-ATM4-SP".

It is interesting to note that there exist around 2.5 K as-set objects that are referenced in the aut-num objects, but are not existing in our collected public IRR datasets. There are three possible main causes (among many) for such aut-num objects; outdated policies in the aut-num objects, typos while registering policies in the aut-num objects, and it is also possible that these as-set objects are registered in the Routing Registries that are not publicly available.

3.2.4 Extracting AS Links and Policies from IRR

We now present our methodology to extract AS links and routing policy annotations from aut-num and as-set objects by referring to a sample routing policy

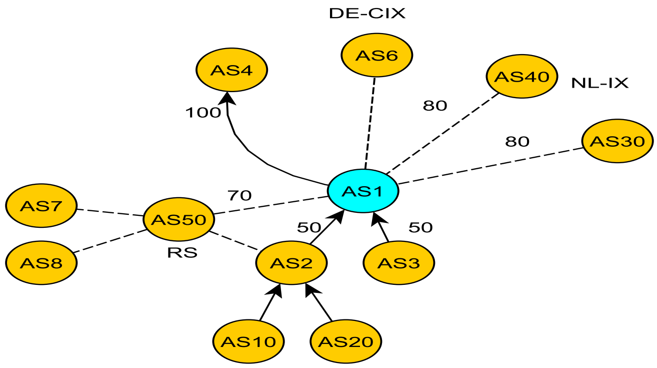

Figure 3.2 AS-level links and annotations extracted from a sample policy of AS1. The lines with solid arrows represent customer-to-provider (c2p) links and dotted lines represent peer-to-peer (p2p) links between ASes with the annotation showing a local preference (LocalPref) value or a peering location (e.g., DE-CIX).

of AS1 in Figure 2.1. Figure 3.2 shows the result of applying our methodology on the AS1's routing policies given in Figure 2.1.

Bilateral Peerings: In lines 13-18, we observe the import and export policies of AS1 towards AS2 and AS3. From lines 13-14, we extract an AS link **AS1-AS2** and can also annotate this link with the local preference (pref [3]) value of 50. We can also extract links **AS2-AS10** and **AS2-AS20** by finding the referrer-AS of the as-set object "AS2:AS-Customers" using our referrer-AS mapping methodology, which is AS2. From lines 15-16, we can only extract the AS link **AS1-AS3** as there is no as-set object referenced in the policy. In lines 17-18, we find an export policy of AS1 towards its customer ASes; the keyword *ANY* specifies that AS1 will announce any prefix announcements to the as-set AS1:AS-Customers, i.e., to the AS links **AS1-AS2** and **AS1-AS3**. Similarly, we extract the link **AS1-AS4** from lines 20-23.

IXP Peerings: If we observe an IP address (v4 or v6) in import or export policies, then we check whether the IP address belongs to an IXP by checking

[3] Preference (pref) is opposite to BGP local preference (LocalPref) in that the smaller values are preferred over larger values.

against the list of IXP prefixes constructed in Section 3.2.1. For example, in line 25, we observe an IP prefix that belongs to DE-CIX. Thus, a link **AS1-AS6** is considered to be an IXP peering link located at DE-CIX. The same methodology has been used by Augustin *et al.* [29] to extract IXP peerings from the IRR. However, as many ASes do not register IP prefix information in their policies, this method can miss a large number of IXP peerings in the IRR. To overcome this limitation, we further look for cases where an as-set object specifies an IXP name (by cross-comparing with those of IXP database built in section 3.2.1, through substring matching). More specifically, we find that ASes register their IXP peers in an as-set object and the name of the as-set object clearly specifies the peering AS as well as the IXP. For example, lines 30-33 show the routing policies of AS1 towards its peers at NL-IX (Netherland Internet Exchange). Thus, we can extract the IXP peering links **AS1-AS30** and **AS1-AS40** by finding the referrer-AS of the as-set object "AS1:AS-PEERS-NLIX" using our as-set to referrer-AS mapping and (IXP name) substring matching.

Multilateral Peerings (MLP): Lines 35-38 show the routing policies of AS1 towards a router server AS (AS50). We look for an as-set object that stores the members of this router server, which is the as-set object "AS50:AS-RS-Peers" (line 9). We extract links **AS1-AS7** and **AS1-AS8**, which are allowed by the import and export policies of AS1 (while AS2 is NOT allowed, according to lines 35-36). It is interesting to note that AS1 is not peering with its customer ASes at the route server, possibly for business reasons. For example, by using a community attribute (line 37), AS1 states that it does not export routes to AS2 through the router server (AS50). The community value {50:50, 50:2} means to allow **ALL** (50:50) but **EXCLUDE** 50:2. This is a common **ALL+EXCLUDE** pattern, used by several IXPs to filter routes that are to be sent to the MLP member ASes [63].

Table 3.1 An example of AS relationship inference.
ASx's routing policy for ASy

Import	Export	Relationship
1. ASy	ASx	Peer-to-Peer (p2p)
2. ASy	ANY	Provider-to-Customer (p2c)
3. ANY	ASx	Customer-to-Provider (c2p)

3.3 AS Relationship Inference

We now present our methodology to infer AS relationships using the routing policy annotations extracted along with AS links.

AS links observed in aut-num objects: Table 3.1 shows the following three policy registration practices of an AS (ASx) that can be used to infer its relationships with a neighboring AS (ASy): (1) ASx does not register the keyword **ANY** in its import and export polices towards ASy, then we classify the link as of type peer-to-peer (p2p). In other words, in a p2p relationship, ASes import only objects (e.g., as-set or route-set objects) maintained by their neighbor ASes and export only objects maintained by themselves. (2) If ASx registers the keyword **ANY** in its export policy for ASy but accept only ASy in its import policy, then we classify the AS link as of type provider-to-customer (p2c), i.e., ASes send all routes to their customer ASes in a p2c relationship. (3) If ASx registers the keyword **ANY** in its import policy from ASy but announce only ASx in its export policy, then we classify the AS link as of type customer-to-provider (c2p); ASes accept all routes from their provider ASes in a c2p relationship. After inferring c2p relationships, similarly to other AS relationship datasets [22, 25], we reverse the direction of the AS link and store it as of type p2c.

s2s relationships: While mapping as-sets to referrer-AS(es) in Section 3.2.3, we find many ASes belonging to the same organization. Such AS-to-organization mapping can be used to generate s2s relationships. However, we do not consider s2s relationships further in the thesis, as we find only a very small fraction

(0.08%) of IRR AS links that are of type s2s. Thus, we only consider p2c and p2p type of AS relationships in the thesis.

AS links observed in as-set objects: Since we do not find routing policy annotations for the AS links only observed in as-set objects, we use the following three as-set objects naming conventions to infer AS relationships: (i) ASes name their as-set objects to specify whether the as-set object is composed of their customer ASes or peer ASes. Thus, AS links from the as-set objects whose name contains texts like "customer", "downstream", or "client" are classified as of type p2c. (ii) ASes name their as-set objects to specify the location of their BGP peerings, e.g., as-set object "AS2:AMS-IX" specifies the peering ASes of AS2 at AMS-IX. Thus, we classify links observed in as-set objects containing abbreviations of IXP names as of type p2p. Most ASes setup p2p relationships at IXPs though other type of relationships are also possible [31]. (iii) ASes name their as-set objects with a text like "upstream" to specify their provider ASes. Thus, links from such as-set objects are classified as of type c2p.

For as-set objects with no hints (in the name) about any AS relationships, if the exporter-AS (or referrer-AS) of an as-set object exists in the as-set object as a member AS, then the as-set object consists of customer ASes of the exporter-AS (or referrer-AS) . Consequently, all the AS links in the as-set object are classified as of type p2c. More specifically, due to similar routing policies for customer ASes and its own AS, ASes often register their own ASes as a member AS in an as-set object containing customer ASes. However, since policies can be different for peers and provider ASes, ASes do not register their own AS as a member AS in an as-set object containing its peer or provider ASes.

Chapter 4

Datasets

In this chapter, we briefly describe the AS topology extracted from IRR and other AS topologies: BGP based (IRL [60]), IXP cliques, and traceroute based (Ark [61], iPlane [40]), all of which are used in this thesis. Then, we also describe the AS relationship datasets which are used to evaluate our proposed AS relationship inference algorithm (in Chapter 6).

4.1 AS Topologies

IRR-based: Using our methodology in Section 3, we extracted 646,431 AS links from the IRR. We filtered out the following AS links: (i) 3,331 AS links where at least one of the AS is a private one; (ii) 2,532 AS links where at least one of the ASes is a route server AS since these links do not carry data traffic; (iii) 30,041 AS links that have not been visible in BGP for over a year now. After filtering out the above AS links, we finally obtain the AS topology snapshot of 610,527 AS links, which is referred to as IRR in Table 5.2. We then infer and classify AS relationships (using our methodology in Section 3.3) into the following: IRRp2p (389,451 AS links), IRRp2c (220,556 AS links), and IRRs2s (520 AS links). Notice that the portion of s2s AS links is negligible

(0.08%), which is skipped.

BGP-based: We combined AS topology datasets shared by IRL [60], Isolario [75], and Cyclops [24], which are largely overlapping due to the shared BGP traces from RouteViews [49], RIPE-RIS [65], and Packet Clearing House (PCH) [66]. However, each of these datasets is using some unique data sources such as Internet2 [67] by IRL and BGPmon [76] by Isolario. The combined AS topology from the above datasets contains 218,319 AS links (and 47,169 ASes). By combining the AS links from LG servers (as described in Section 3.1) and other BGP-based datasets, we obtain an AS topology snapshot of 239,037 AS links (and 48,097 ASes), which is referred to as **BGPAll** hereafter. Overall, we discovered 20,718 unique AS links by querying LG servers, i.e., around 8.6% addition to the BGP-based AS topology observed in IRL, Isolario, and Cyclops.

IXP Peerings: As explained in Section 3.2.1, we collected a list of IXP participants from Isolario (221 IXPs with their member ASes) and PeeringDB (379 IXPs with their member ASes). Moreover, since ASes can also connect to ASes in private peering facilities, we also collect a list of 1,190 private peering facilities from PeeringDB [62]. For example, Telehouse London (Docklands North) is a private peering facility with 253 member ASes. Since the IXP peering matrices are not publicly known [21, 26, 31], it is difficult to figure out how many peering links are operational at IXPs. Thus, as similar to He *et al.* [21], we create a superset of all possible IXP links; we assume that the participants of an IXP form a clique. By combining the cliques of all IXPs shown in our dataset, we come up with the dataset of 965,461 AS links which contain 6,596 ASes. We refer to this dataset as **IXPAll**.

Traceroute-based: We obtained 104,809 AS links from the CAIDA Ark dataset [61] and 43,443 ones from the iPlane [40] dataset during Oct. 2013. A combined view of the two datasets is referred to as **TrAll**, which contains 115,879 AS links (and 40,216 ASes). Note that we use only the recently (and regularly) published AS topology datasets, not including ones such as Ono [26], which had been collected using BitTorrent P2P clients in 2007-2008 due to staleness concern. Likewise, we exclude DIMES [39] datasets as they had not been updated since Apr. 2012.

Combined AS Topology: In total, we combined all the AS links observed in BGPAll, IXPAll, and TrAll, which consists of 1,110,403 AS links (and 48,790 ASes). We refer to the combined dataset as **CombinedAll**.

4.2 AS Relationship Datasets

The following AS relationship datasets are used to compare with our proposed AS relationship inference method. Note that the results of AS relationship algorithms are based on AS paths observed in BGP in the month of Oct. 2013, unless otherwise specified.

- **Gao**: We ran scripts shared by Gao [9] on the BGP AS paths extracted from RouteViews BGP traces to infer AS relationships. We find 92,143 p2c links and 2,553 p2p links in Gao.

- **Cyclops**: Oliveira *et al.* [24] proposed an algorithm to infer AS relationships, which begins from a set of Tier-1 ASes (which are listed on Wikipedia [69]) and infers p2c relationships for links observed in these ASes. Note that all other remaining links are of type p2p. We find 93,210 p2c links and 34,663 p2p links in Cyclops.

- **Isolario**: Gregori *et al.* [25] used a very similar approach as Cyclops; for each AS path, their algorithm identifies possible relationships and then infers the actual relationship based on lifetimes of the AS paths. We find 100,882 p2c links and 89,732 p2p links in Isolario.

- **CAIDA AS Relationships**: Luckie *et al.* [22] refined existing AS relationship inference methods that are based on AS paths in BGP, and validated a large number of inferred AS relationships by collecting ground truth information (i) directly reported by network operators, (ii) extracted from the IRR RPSL objects, and (iii) obtained from BGP community values in BGP traces. We find 86,739 p2c links and 66,617 p2p links, with this method.

- **GT-RPSL**: This is the ground truth dataset shared by Luckie *et al.* [22]. They extracted 6,530 p2c relationships from routing policies registered in the RIPE IRR dataset of Apr. 2012.

- **GT-Comm**: This is another ground truth dataset share by Luckie *et al.* [22]. They extracted 41,604 relationships (16,248 p2p and 23,356 p2c) by using a dictionary of 1,286 BGP community values from 224 ASes, which is constructed from the BGP traces of Apr. 2012.

Chapter 5

Analysis

In this chapter, we first present a comparison of ASes sharing their BGP traces publicly (BGP feeders) with the ASes who we have queried to collect LG servers. Second, we analyze the current practice of ASes registering their routing policies in the IRR. Finally, we analyze AS topology datasets to compare the AS links observed in IRR with the ones observed in BGP, traceroute, and IXP cliques. Note that the datasets of BGP, traceroute, and IXP cliques were collected during Oct. 2013.

5.1 Comparison of BGP feeders

There are three popular BGP collector projects: RouteViews [49], RIPE-RIS [65], and PCH [66]. The ASes sharing their BGP traces to the BGP collector projects are known as BGP feeders [38]. In this section, we are interested in finding out whether, by querying LG servers, we can collect BGP traces from ASes that are not BGP feeders of RouteViews, RIPE-RIS, and PCH. Such analysis indicates whether BGP traces collected from the LG servers help discover new AS links that are not found in the other AS topology datasets (e.g., IRL [60]). Moreover, collecting BGP traces from new BGP feeders help widen our limited view of

BGP observed from current BGP collectors [38].

We have collected information regarding the BGP feeders (i.e., ASNs and IP addresses of routers) of RouteViews [49], RIPE-RIS [65], and PCH [66] from their websites in the month of March 2013. The comparison between the BGP feeders of different projects are based on the ASN and IP address of BGP routers sharing the BGP traces. That is, if the ASN and/or IP address of a BGP router matches between the LG servers and RouteViews BGP feeders, then it is considered that the same BGP router (of an AS) is sharing its BGP traces with both RouteViews and the LG servers.

Table 5.1 shows the number of common BGP feeders (ASes and routers) sharing their BGP traces with the RouteViews, RIPE-RIS, PCH, or LG server datasets. The diagonal (in bold) is the number of BGP feeders available only in one dataset; either in RV, RIPE, PCH, or LG servers. We observe differences in the number of ASes and router IPs overlapping between different datasets. For example, 63 neighboring ASes of LG servers are sharing their traces with RV. However, only 36 router IPs are matched between LG servers and RV. Further investigation leads us to find the following two reasons for such mismatches: (i) An AS can be peering on an IPv4 connection with RV while on an IPv6 connection with LG servers. In that case, when two datasets are compared to check for the overlapping ASNs and router IPs, the observed router IPs can be different in both datasets though they are with the same ASN. (ii) An AS can be peering with RV at a different location in the Internet from where an LG server is located, thus the observed router IPs between the two datasets can be different as well, while they have the same ASNs.

Moreover, we find that 545 (out of 1.1 K) neighboring ASes of the LG servers overlap with RouteViews, RIPE-RIS, or PCH. More importantly, we observe that 620 neighboring ASes of the LG routers are not sharing their BGP traces with RouteViews, RIPE-RIS, nor PCH. We further inspect the number of routes announced by each neighboring ASes of LG servers to find that 70%

(of the 1.1 K) neighboring ASes of the LG servers announce a small number (1 to 100) of BGP routes, since most of these ASes are stub ASes. The remaining (30%) neighbors of LG servers announce BGP routes in the range of 100 to 450 K. Overall, we were able to collect 128 BGP routing tables of around 450 K prefixes from the LG servers, which is approximately equal to the size of full BGP routing table in the current Internet [53].

The analysis presented so far in this section suggests that there are many ASes who are willing to publicly share their BGP traces by operating LG servers, which in turn begs the question that why such ASes have yet to offer feeds to route collectors. We suggest two possible reasons: (i) In the past, network operators were motivated to share their BGP feeds to the route collectors in order to advertise their rich connectivity and dominance (especially Tier-1's) in the Internet [38]. However, they may not need to do that any more as maintaining an LG server serves that purpose too. Besides, maintaining an LG server by an AS is helpful for operational reasons such as troubleshooting routing issues. (ii) BGP collector projects such as RouteViews have presence at a limited number of locations in the Internet (e.g., large IXPs) and mostly collect traces from ASes present at those locations (e.g., members of large IXPs). Thus, RouteViews can not collect traces from the ASes which are not located at these locations but are sharing their feeds to the LG servers.

Table 5.1 The number of overlapping and unique (in bold) ASes and peering routers between various BGP feeders (RouteViews (RV), RIPE-RIS, PCH, and LG servers).

Collector (Total # of ASes and Routers)	RV ASes (Routers)	RIPE	PCH	LG servers
RV (179 and 368)	**72 (276)**	46 (27)	76 (44)	63 (36)
RIPE (343 and 599)	46 (27)	**51 (314)**	235 (215)	191 (133)
PCH (1.2 K and 2.7 K)	76 (44)	235 (215)	**719 (2 K)**	428 (615)
LG servers (1.1 K and 3.3 K)	63 (36)	191 (133)	428 (615)	**620 (2.6 K)**

5.2 Registration of Routing Policies in the IRR

In this section, we analyze the current practice of ASes registering their routing policies in the IRR.

5.2.1 Policies in aut-num Objects

To find currently operational ASes, we first collected the RIR dataset (as of Nov. 1st, 2013). From the dataset, we observed total 56,718 ASNs; they belong to one of the following regions: RIPE NCC (40%), ARIN (38%), APNIC (14%), LACNIC (6%), and AfriNIC (2%). We find 26,657 (47% of 56,718 allocated ASNs) are registering their routing policies in the IRR aut-num objects. Figure 5.1 shows the fractions of ASes (for different RIR regions) that register their policies in terms of the *aut-num* and *route objects*, respectively. We find that route objects are usually registered across the RIRs. This is because ASes may want to publicly announce the ownership of their IP prefixes, which can then be used for IP trading or troubleshooting purposes for Internet routing [5]. Also, ASes may be asked by their providers or peer ASes to register their route objects in the IRR to ensure the routability of their prefixes [7]. Note that registering aut-num objects is more prevalent in RIPE (92%) and APNIC (50%) than in other RIRs.

We further investigate how ASes register their routing policies in aut-num objects. Figure 5.2 shows the three most widely used combinations of routing policy entries registered in the period starting from Oct. 2010 to Nov. 2013. We find that a large number of ASes register their import and export policies (i.e., import_export) in the IRR, and the trend is consistent across the observed period. This indicates that we can infer AS relationships of a large portion of the routing policies registered in the IRR, as explained in Section 3.3. Note that only a small number of ASes register their import (or export) policies only, which cannot help infer AS relationships.

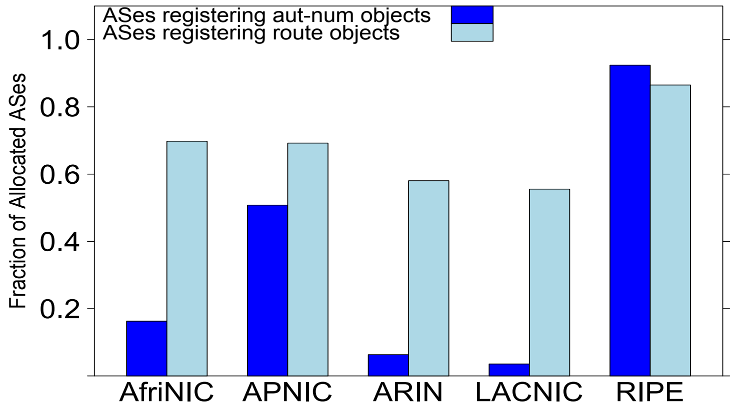

Figure 5.1 Policy registration practice of ASes in IRR.

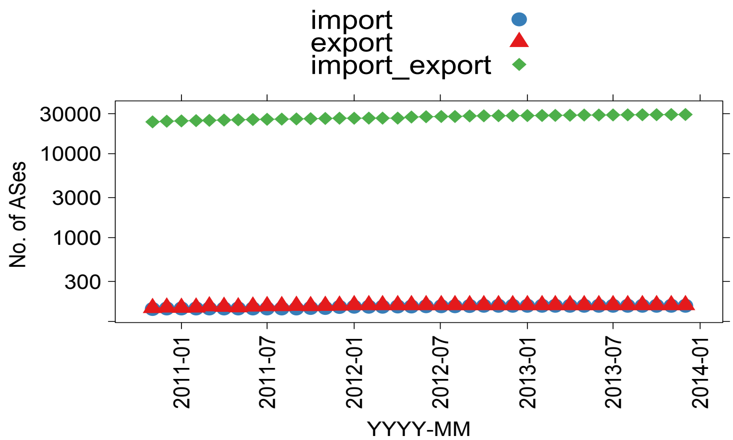

Figure 5.2 Import/export routing policies in aut-num objects are registered differently per AS.

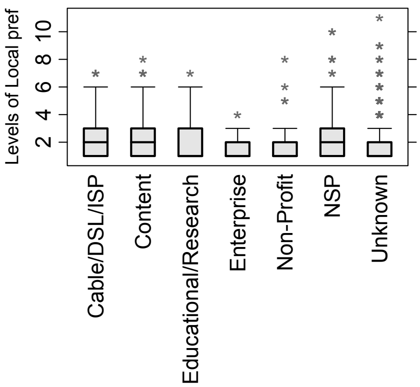

(a) AS business types

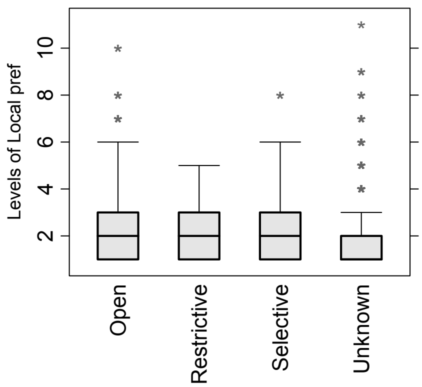

(b) Peering policies

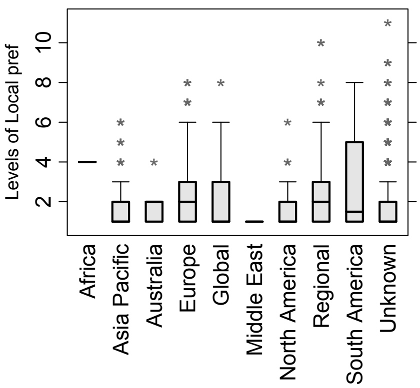

(c) Operational scope of ASes

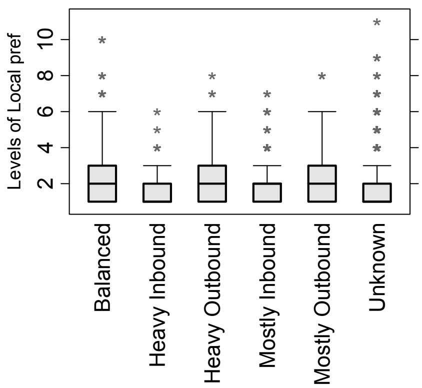

(d) Traffic types

Figure 5.3 The distributions of the number of Local preference (LocalPref) values registered by individual ASes in the IRR are shown based on PeeringDB [62].

5.2.2 Number of Local Preference (LocalPref) Values in the IRR

The local preference (LocalPref) attribute is one of the most important "knobs" for BGP routing. We observe 16,619 ASes that register their LocalPref values in the IRR; 15,124 ASes register only a single LocalPref value while 1,495 ones register 2 to 11 values. For example, *vk.com* (AS47541 located in Europe) is the second largest social networking service, and uses 8 different LocalPref values to control traffic. Since LocalPref is non-transitive, it is not transferred to neighboring ASes. Thus, the IRR is a good source for understanding the usage/number of LocalPref values across different ASes.

To further explore the LocalPref registration practice of ASes, we use the following four classifications of 3,626 ASes from the PeeringDB [62] snapshot (as of Nov. 1st, 2013): (1) business types, (2) peering policies, (3) operational scopes, and (4) traffic types. The business types include Cable/DSL/Access Provider (1,224), Network Service Provider (NSP) (1,130), Content Provider (885), Education/Research (140), Enterprise (134), and Non-Profit (111). The peering policies contain open (2,743), selective (772), and restrictive (109). The operational scope of an AS includes Regional (1,301), Europe (1,021), Global (576), North America (327), Asia pacific (305), South America (54), Australia (15), Africa (12), and Middle East(3). For the traffic types, there are Balanced (1,454), Mostly Inbound (946), Mostly Outbound (799), Heavy Outbound (272), and Heavy Inbound (153). We also classify the 3,626 ASes from an RIR perspective: RIPE (1,881), ARIN (897), APNIC (482), LACNIC (243), and AfriNIC (67). While the continental distribution of ASes in PeeringDB is somewhat biased, non-stub ASes are geographically scattered across the entire Internet [32].

Figure 5.3 shows the boxplot of the number of LocalPref values used by each AS in the IRR. Note that ASes found in the IRR but not in PeeringDB belong to the "Unknown" category. Figure 5.3a first shows that ASes whose types are Cable/DSL/ISP, Content, and NSP exhibit higher median LocalPref values (i.e., 2) than others, which implies the practice of traffic load balancing for such ASes. We also show that LocalPref usages are very similar across different types of peering policies in Figure 5.3b; their median value is 2. We observe that regional ASes tend to have greater number of LocalPref values than global ASes as shown in Figure 5.3c. Finally, Figure 5.3d shows that ASes with traffic types of Balanced, Heavy Outbound, and Mostly Outbound register the higher number of LocalPref values in the IRR than other ASes. For example, ASCENT (AS52925, NSP, open, South America, Balanced) provides in a range of services such as cloud computing, managed hosting, and co-location technologies, and hence uses eight LocalPref values for load balancing traffic towards its providers and peer ASes.

5.3 Analysis on AS-level Topology

In this section, we analyze AS topology datasets to compare the AS links observed in IRR with the ones observed in BGP, traceroute, and IXP cliques. Note that the datasets of BGP, traceroute, and IXP cliques were collected during Oct. 2013.

Table 5.2 The portion of IRR AS links that are overlapping with other datasets is shown.

Name	# of Links	BGPAll	TrAll	IXPAll	CombinedAll
IRR	610,527	24%	10%	57%	68%
IRRp2c	220,556	30%	20%	16%	40%
IRRp2p	389,451	17%	1%	75%	76%

5.3.1 Overlapping and Missing IRR-based AS Links

Analysis of overlapping and missing links

We first compare IRR-based AS links with the other datasets to find overlapping or unique AS links among them. Such analysis is important in quantifying how many AS links are newly extracted from the IRR or which AS links are missing in the IRR. Table 5.2 shows the fractions of IRR-based AS links which are overlapping with the other datasets. We find that a small fraction of IRR links are overlapping with the BGPAll (24%) and TrAll (10%). However, a large fraction (57%) of IRR links are overlapping with the IXPAll, which implies the practice of widespread IRR registrations for the member ASes of IXPs. Most of IRR-based AS links which are overlapping with IXPAll are of p2p type (75%). However, we find there are also AS links of p2c type (16%), which means that the ASes may connect with their customers through IXPs. Note that only 1% of the p2p type of IRR-based AS links are overlapping with TrAll, which is possibly due to the limited vantage points (i.e., measurement locations) in iPlane [40] and CAIDA Ark [61]. Overall, 195 K (32%) IRR AS links are not found in any datasets. This means a large fraction of both p2p and p2c types of links are missing. A possible reason for missing p2p links is that Isolario only publishes the list of member ASes of 221 IXPs, while there are more than 400 IXPs operating in the Internet. On the other hand, a less frequently updated list of the member ASes of 379 IXPs is available in PeeringDB.

We next investigate why the other datasets often exclude AS links from IRR. The well-known main reason is the stale information in the IRR. To verify this, we plot the fraction of missing IRR-based links per AS in BGPAll across different regions in Figure 5.4. Since a large number of AS links extracted from most recently updated routing polices are also missing across different datasets, we conclude that the stale information of routing policy is not the only reason for the missing IRR-based AS links. There are two further possible reasons for missing IRR-based links. First, the incompleteness of the BGP-based dataset has been reported in the literature [24,36,38], which is due to the limited number of ASes sharing their BGP feeds with RouteViews and RIPE-RIS. Moreover, an AS topology view from the BGP-based dataset may also be biased, as current route collector projects have better views of the core of the Internet rather than the other parts; tier-1 ISPs more actively share their BGP traces than other ASes [38]. Second, the traceroute-based dataset suffer from limited vantage points, selectively probing prefixes, IP-to-AS mapping issues [37], and inability to discover backup links [40].

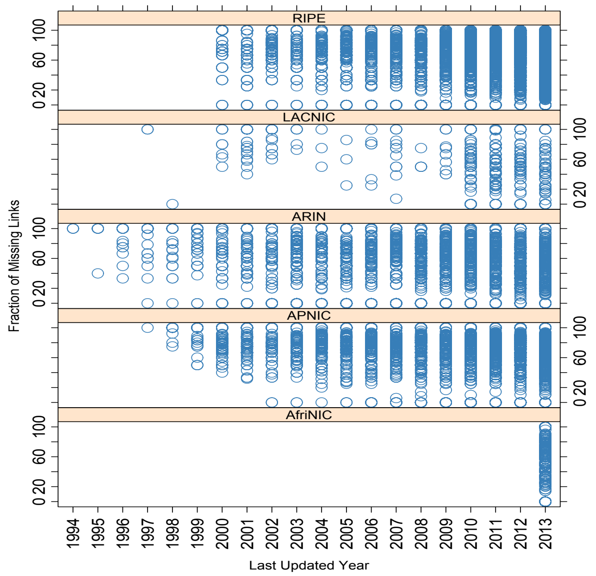

Figure 5.4 Fraction of per-AS IRR Links missing in BGPAll is shown depending on regions and IRR update year.

Practice of registering AS links

To investigate how AS links are registered in the IRR, we first classify IRR-based AS links into the following four types:

(1) **InAutnumDual** indicates that both of ASes that share the same AS link register their aut-num objects. Such links are often assumed to be more reliable since both ASes register the existence of the link in IRR [21, 22].

(2) **InAutnum** indicates that only one AS of a given AS link registers its aut-num object.

(3) **InAutnumThirdParty** includes an AS link identified through the aut-num object of an AS that does not belong to the AS link. This often happens

when an AS refers to an as-set object, which in turn refers to other as-set objects (sub as-set objects). This type often includes the ASes which register their as-set objects but do not register their policies in aut-num objects.
(4) **InASSet** specifies AS links identified only in the as-set objects. As shown in Section 3.2.3, many as-set objects are not referenced in any of the aut-num objects as some ASes are not interested in maintaining their routing policies in aut-num objects.

We further divide the above types into subtypes with respect to p2c and p2p relationships. For example, InAutnumDual have two subtypes: **InAutnumDualp2c** is for p2c links and **InAutnumDualp2p** is for p2p links.

Table 5.3 shows how many IRR-based AS links (of each type) are overlapping with BGPAll, IXPAll, and CombinedAll, respectively. First, we investigate how ASes register their routing policies in IRR by looking at the number of IRR-based links of different types. We find both of ASes in 157 K AS links (26% of total IRR-based links) are registered, and only one AS in 325 K IRR AS links (53%) are registered. The remaining IRR-based AS links are either extracted from the third party routing policies (7%) or from the as-set objects (14%) that are not referenced in any of the aut-num objects. Second, we find that a relatively similar portion of IRR-based links are overlapping with BGPAll across the 4 types, which means that not only InAutnumDual, but the other 3 types are also reliable IRR registration practices. Third, similarly, we observe 94.5% of AS links of InAutnumDual are overlapping with CombinedALL. However, 65.9% of AS links of InAutnum also match with CombinedAll. Finally, we observe that the p2c type of links are most missing.

Overall, we conclude that registering their policies in the IRR by ASes are largely driven by their own business or operational needs since a substantial portion of non-InAutnumDual type AS links are overlapping with CombinedAll. We also find that many ASes may update their policies in the IRR less frequently than actual changes observed in BGP. Another negative implication is that

ASes can suddenly stop registering their detailed routing policies in the IRR. For example, we observe that Hurricane Electric (AS6695) had been registering a very detailed routing policies until Apr. 2013, and then it has removed them and has started registering "From AS-ANY accept ANY", perhaps to show their open peering policies.

Table 5.3 IRR AS links overlapping in BGPAll, IXPAll, and CombinedAll datasets.

Name	# of Links	BGPAll	IXPAll	CombinedAll
InAutnumDual	157,090	25.2%	90.97%	94.5%
InAutnum	324,996	20.9%	58.1%	65.9%
InAutnumThirdParty	45,485	26.7%	10.8%	35.3%
InASSet	82,956	34.1%	14.8%	44.2%
InAutnumDualp2c	14,290	59.8%	54.1%	86.2%
InAutnump2c	78,633	36.9%	21.9%	50.9%
InAutnumThirdPartyp2c	45,219	26.7%	10.8%	35.3%
InASSetp2c	82,414	33.9%	14.6%	44%
InAutnumDualp2p	142,703	21.7%	94.7%	95.4%
InAutnump2p	245,940	15.8%	69.8%	70.7%
InAutnumThirdPartyp2p	266	24.4%	22.9%	31.2%
InASSetp2p	542	63.5%	47.6%	73.1%

Usages of IRR by member ASes of IXPs

We next investigate usages of the IRR by member ASes of IXPs, which can help in detecting a large number of IXP peerings. We observe that IRR registration practices of member ASes of IXPs are diverse. Member ASes of 105 IXPs register IP prefixes of their IXP peerings in their policies in aut-num objects. On the other hand, member ASes of other 110 IXPs manage their IXP peerings through as-set objects and specify their multilateral peerings with a route server by specifying either the route server ASN or the as-set object of route server. We observe that only 226 ASes use BGP community values to filter their multilateral peerings and that is to restrict their route announcements to their customers and private peers through a route server.

Table 5.4 shows top 15 IXPs in the Internet in terms of the number of IXP member ASes whose information are from Isolario [75]. We observe that AMS-IX and DE-CIX, the two largest IXPs in the Internet, have a large fraction of their peering matrices registered in IRR. This is also verified by the website of AMS-IX [63], which encourage its member ASes to use IRR instead of BGP community values for the purpose of filtering. However, since some IXPs rarely use IRR, we cannot find their peering matrices in the IRR. For example, member ASes of PTTMetro-SaoPaulo does not actively use IRR.

Giotsas *et al.* [31] inferred multilateral peerings of 13 European IXPs using the BGP community values observed in BGP traces; their datasets are not publicly available so we cannot directly compare with it in this thesis. However, we find that peering densities for DE-CIX (79%) and MSK-IX (95%) as we observed in IRR are same with the peering densities reported in their work. On the other hand, 95% peering density is reported for PLIX in their work but we find 7.8% peering density in IRR, which is due to the less active usage of the IRR by PLIX member ASes.

5.3.2 BGP-based AS Links vs. IRR-based AS Links

Since BGP-based AS topology is often considered as more accurate than other sources, we compare AS links in BGPAll with IRR-based ones to investigate how many BGP links are also observed in IRR. Overall, we observe that 62% of 239,037 AS links in BGPAll are also observed in IRR.

To further explore how different RIR regions use IRR, we show the fraction of per-AS links in BGPAll which are also observed in IRR in Figure 5.5. While LACNIC and AfriNIC regions show less active usage of IRR, the missing BGP links are observed in all the RIR regions. Moreover, we find that a large number of ASes are registering their partial AS links in the IRR. One of the reasons for

Table 5.4 Top 15 IXPs (in terms of member ASes reported by Isolario) and their peering matrices in IRR and a Combined (BGP, Traceroute, and IRR) dataset.

IXP	**Region**	**ASes**	**Clique**	**IRR**	**Combined**
AMS-IX	EU	607	183,921	60.81%	64.20%
DE-CIX	EU	507	128,271	78.39%	81.15%
LINX	EU	505	127,260	45.83%	52.47%
PTTMetro-Sao Paulo	LA	445	98,790	2.47%	11.88%
MSK-IX	EU	374	69,751	95.26%	95.38%
NL-IX	EU	341	57,970	20.39%	22.35%
FranceIX	EU	240	28,680	21.27%	29.84%
PLIX	EU	227	25,651	7.82%	12.44%
Any2-CA	NA	221	24,310	6.63%	18.27%
SIX	NA	188	17,578	7.25%	15.23%
DATAIX	EU	186	17,205	37.79%	40.61%
HKIX	AP	177	15,576	9.70%	15.68%
TorIX	NA	172	14,706	6.21%	9.98%
SwissIX	EU	143	10,153	60.28%	64.33%
KleyReX	EU	137	9,316	28.99%	31.03%

this partial registration is the existence of non-disclosure peering agreements [4]. Since ASes are not allowed to disclose their peering relationships, they do not register these AS links in IRR.

Finally, we observe that some large ASes show a significant usage of the IRR. For example, RETN (AS9002) in RIPE region has 91% of its 2,048 links in BGP, which are also observed in IRR. However, many other large ASes do not actively use IRR. For example, Verizon (AS701) and AT&T (AS7018) show very poor usage of IRR, i.e., only 3% of their 2 K links in BGP are observed in IRR. We also notice that both Verizon and AT&T have not presented at IXPs. Thus, their links are neither reported by themselves in IRR nor by their peering ASes.

5.3.3 AS Degree Distribution

In this section, we investigate whether the different methods of collecting AS topologies result in different AS degree distributions. The **AS degree distribution** is the probability that a randomly selected AS is k-degree: $P(k) = n(k)/n$; where n is the number of ASes and $n(k)$ is the number of ASes with degree k. The degree distribution is the most frequently used topology characteristic [20].

Figure 5.6 shows the PDF of the AS topology datasets plotted along with that of the IRR derived AS links. We observe that IRR has many more AS links for the moderate degree ASes which implies that the IRR is popular in the realm of smaller ISPs. Large ISPs have shown little interest in the IRR as it is difficult to manage complex routing policies in the IRR [6]. We also observe that traceroute-based AS topology (TRAll) closely match for the low degree ASes. However, traceroute reports more AS links for higher degree ASes. To find out the reasons, we analyze the traceroute IP paths collected from the iPlane in March 2013. We observe that iPlane has a selective list of IP prefixes (120 K out of approximately 450 K IP prefixes that are currently operational in BGP [53]) to probe the Internet and this list seems to concentrate more

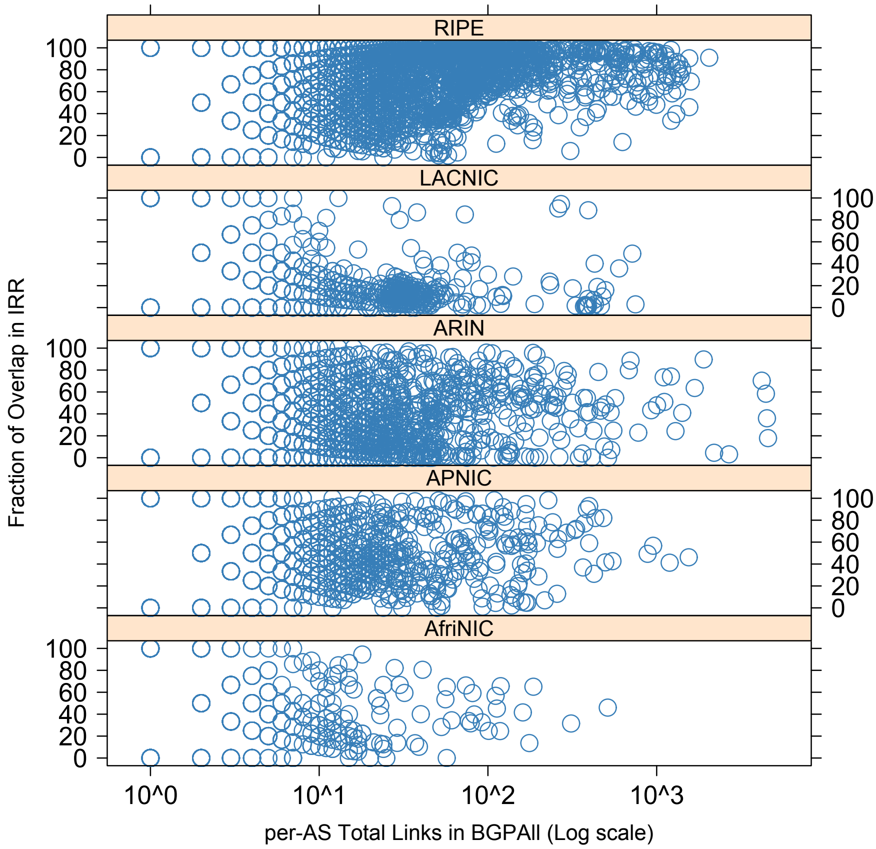

Figure 5.5 per-AS BGP Links overlapping in IRR.

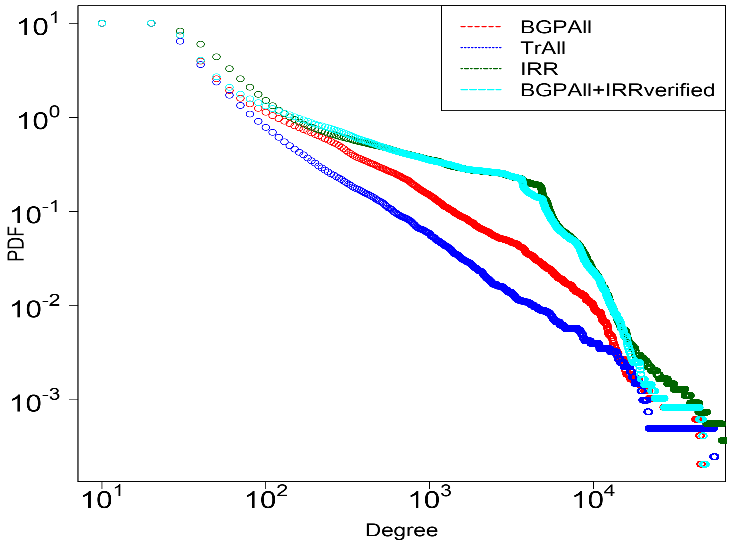

Figure 5.6 AS Cumulative degree distribution (PDF) of the AS topologies extracted from BGP (BGPAll), traceroute (TrAll), IRR and verified one (BG-PAll+IRRverified).

on probing the core Internet [40]. Thus, iPlane discovers more connectivity of higher degree ASes. Overall, traceroute based projects such as Ark and iPlane suffer from limited vantage points, selectively probing IP prefixes, errors in the translation from IP to AS Path. Such factors impact the AS topology view observed from traceroute measurements, which can also be error prone [16,37].

Chapter 6

AS Relationship Inference

In this chapter, we evaluate our proposed AS relationship inference method (described in Section 3.3) against existing methods that are based on BGP AS paths. We also compare our results with two ground-truth datasets shared by Luckie *et al.* [22].

6.1 Evaluation Results

As shown in section 5.3, we infer 389,451 p2p and 220,556 p2c AS relationships from the IRR. Table 6.1 shows the fraction of the inferred AS relationships matching with those of the other existing algorithms. We observe that a high fraction (90%) of p2c relationships are consistently matched with other datasets. However, we find significant improvements of Isolario and CAIDA methods over Gao and Cyclops ones in the inferences of p2p relationships, which has been achieved by relying less on the valley-free property; some AS paths in the BGP do not follow the valley-free property due to BGP mis-configurations, poisoned paths, or special routing policies [22].

Interestingly, 99% of our inference results are matched with GT-RPSL, which is extracted from the IRR by evaluating the policies of both-end ASes

Table 6.1 IRR AS Relationships compared with others.

Name	Matching p2c (accuracy%)	Matching p2p (accuracy%)
Gao	57,530 (92.5%)	1,180 (81.2%)
Cyclops	53,196 (89.2%)	15,096 (58.1%)
Isolario	61,651 (90.8%)	48,871 (92.8%)
CAIDA	57,569 (96.6%)	37,184 (94.7%)
GT-RPSL	6,253 (99.0%)	
GT-Comm	15,412 (96.4%)	10,770 (94.1%)

in an AS link. We further highlight that inferring AS relationships from the policy of a single AS in an AS link is also highly accurate. For the mismatching 1% relationships (i.e., 63 AS links), we find that all of the reported p2c relationships in GT-RPSL has recently been changed from p2c to p2p, which also could have been correctly inferred by our inference method, had it not been changed since Apr. 2012. However, we observe that a small number of AS relationships have been changed because we analyze the IRR dataset of Nov. 1st, 2013, but the GT-RPSL dataset was constructed from RIPE IRR of Apr. 2012. We find that the GT-RPSL does not contain any p2p links. On the other hand, we also show that p2p links inferred by our method are also matched well with the ones inferred by other methods. Finally, we find that the accuracy of our methods against the ground-truth dataset extracted from BGP communities (GT-Comm) is also significantly high; 96% for p2c and 94% for p2p. In summary, these performance results imply that inferring AS relationships from the IRR is significantly reliable and promising.

Chapter 7

Summary & Future Work

In this thesis, we first highlighted the less-known capabilities of Looking glass (LG) servers to construct Internet AS topology. By collecting **show ip bgp summary** command responses from 245 LG servers (from 1.9 K locations in 110 countries) and **BGP neighbor ip advertised routes** command responses from 59 LG servers (from 250 locations in 40 countries) in Oct. 2013, we build an AS topology estimate of around 143 K AS links. We newly discovered 20 K AS links and 686 ASes that are not found in BGP, traceroute, and IRR based AS topologies. Clearly, LG servers help in augmenting the current AS topology collection efforts reliably as BGP based methods are less error prone as compared to traceroute-based ones. However, the AS topology view from the LG servers suffers from limited vantage points of the LG servers and BGP export policies employed by the neighboring ASes of LG servers.

Thus, to overcome the limitations of an AS-level view observed through BGP data sources, we presented a methodology to extract an AS-level topology from the IRR. By using our methodology, we extracted 610 K AS links from the IRR; 68% of them are matched in BGP, traceroute, and the cliques of Internet eXchange points (IXPs). We observed an active usage of the IRR by member

ASes of IXPs, which results in inferring peering matrices of many large and small IXPs. Finally, we proposed a method to infer business relationships (e.g., p2c or p2p) between ASes using routing polices stored in the IRR. The overall accuracy of our method is comparable (97% for p2c, 95% for p2p links) to the existing methods that infer AS relationships using only BGP AS paths.

Overall, we conclude that the IRR is a strong complementary source to provide better understandings of the structure, performance, dynamics, and evolution of the Internet since it is actively used by a large number of operational ASes in the Internet. Further analysis and validation of our AS level topology (of 610 K AS links) is needed. For example, to find out how such a comprehensive AS-level topology helps in detecting AS path spoofing attacks in BGP [23]. We envision that more LG servers are going to be deployed and IRR registration practices of ASes will improve further in the future. Thus, the research and operational network community needs to be aware of the facilities provided by them to discover Internet AS topology.

Bibliography

[1] Y. Rekhter, T. Li, and S. Hares. A Border Gateway Protocol. *RFC 4271*, Jan., 2006.

[2] Internet Routing Registry. `http://www.irr.net`.

[3] L. Blunk, J. Damas, F. Parent, and A. Robachevsky. Routing Policy Specification Language. *RFC 4012*, Mar., 2005.

[4] R.A. Steenbergen, A Complete Guide to Peering. *NANOG 51*, Jan., 2011.

[5] D. McPherson, S. Amante, E. Osterweil, and L. Blunk. IRR & Routing Policy Configuration Considerations. (draft-ietf-grow-irr-routing-policy-considerations-02), Nov., 2013.

[6] ISP Route Filtering: Responsibilities and Technical Challenges. *Nanog43*, Jun., 2008. `https://www.nanog.org/meetings/abstract?id=82`.

[7] NTT Communications. `http://www.us.ntt.net/support/policy/routing.cfm`.

[8] M. Caesar and J. Rexford. BGP routing policies in ISP networks, *IEEE Network*, Dec., 2005.

[9] L. Gao. On Inferring Autonomous System Relationships in the Internet. *IEEE/ACM Transactions on Networking (TON)*, Vol. 9, no. 6, 2001.

[10] G. Siganos and M. Faloutsos. Analyzing BGP Policies: Methodology and Tool. In *IEEE INFOCOM*, Mar., 2004.

[11] G. D. Battista, T. Refice, and M. Rimondini. How to extract BGP peering information from the Internet routing registry. *ACM MineNet*, 2006.

[12] H. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger. Towards Capturing Representative AS-Level Internet Topologies. *Computer Networks*, vol. 44, no. 6, April, 2004.

[13] John S. Otto, M. Sanchez, D. Choffnes, F. Bustamante, and G. Siganos. On Blind Mice and the Elephant – Understanding the Network Impact of a Large Distributed System. *ACM SIGCOMM*, Aug., 2011.

[14] X. Dimitropoulos, D. Krioukov, M. Fomenkov, B. Huffaker, Y. Hyun, K. claffy, and G. Riley. AS relationships:Inference and validation. *ACM SIGCOMM Computer Communication Review (CCR)*, vol. 37, no. 1, Jan., 2007.

[15] P. Mahadevan, C. Hubble, D. Krioukov, B. Huffaker, and A. Vahdat. Orbis: rescaling degree correlations to generate annotated internet topologies. *ACM SIGCOMM*, Aug., 2007.

[16] P. Marchetta, W. Donato, and A. Pescapé. Detecting Third-party Addresses in Traceroute Traces with IP Timestamp Option. *PAM*, Mar., 2013.

[17] H. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. Anderson, A. Krishnamurthy, and A. Venkataramani. iPlane: An Information Plane for Distributed Services. In *USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, Nov. 2006.

[18] Y. Zhang, R. Oliveira, Y. Wang, S. Su, B. Zhang, J. Bi, H. Zhang, and L. Zhang. A Framework to Quantify the Pitfalls of Using Traceroute in AS-

Level Topology Measurement. *IEEE Journal on Selected Areas in Communications (JSAC)*, vol. 29, no. 9, Oct. 2011.

[19] A. Khan, H. Kim, T. Kwon, and Y. Choi. Public Internet Routing Registries (IRR) Evolution, In Proc. of International Conference on Future Internet Technologies *(CFI)*, Jun. 2010.

[20] P. Mahadevan, D. Krioukov, M. Fomenkov, X. Dimitropoulos, K. Claffy, and A. Vahdat. The internet AS-level topology: three data sources and one definitive metric. *ACM SIGCOMM Computer Communication Review (CCR)*, vol. 36, no. 1, Jan., 2006.

[21] Y. He, G. Siganos, and M. Faloutsos. Lord of the Links: A Framework for Discovering Missing Links in the Internet Topology. *IEEE/ACM Transactions on Networking (TON)*, vol. 17, no. 2, Apr., 2009.

[22] M. Luckie, B. Huffaker, A. Dhamdhere, V. Giotsas, and k claffy. AS Relationships, Customer Cones, and Validation. In *ACM Internet Measurement Conference (IMC)*, Oct., 2013.

[23] K. Butler, T. R. Farley, P. McDaniel, and J. Rexford. A Survey of BGP Security Issues and Solutions. *Proc. of the IEEE*, vol. 98, no. 1, Jan., 2010.

[24] R. Oliveira, D. Pei, W. Willinger, B. Zhang, and L. Zhang. The (In) Completeness of the Observed Internet AS-level Structure. *IEEE/ACM Transactions on Networking (TON)*, vol. 18, no. 1, Feb., 2010.

[25] E. Gregori, A. Improta, L. Lenzini, L. Rossi, and L. Sani. BGP and Inter-AS Economic Relationships. In *IFIP Networking*, 2011.

[26] K. Chen, D. Choffnes, R. Potharaju, Y. Chen, F. Bustamante, D. Pei, and Y. Zhao. Where the Sidewalk Ends: Extending the Internet AS Graph Using Traceroutes From P2P. *ACM CoNEXT*, Dec., 2009.

[27] A. Khan, H. Kim, T. Kwon, and Y. Choi. A comparative study on IP prefixes and their origin Ases in BGP and the IRR. *ACM SIGCOMM Computer Communication Review (CCR)*, vol. 43, no. 3, Jul., 2013.

[28] A. Khan, H. Kim, T. Kwon, and Y. Choi. Public Internet Routing Registries (IRR) Evolution, In Proc. of International Conference on Future Internet Technologies *(CFI)*, Jun. 2010.

[29] B. Augustin, B. Krishnamurthy, and W. Willinger. IXPs: mapped?. *ACM Internet Measurement Conference (IMC)*, Nov., 2009.

[30] B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger. Anatomy of a Large European IXP. *ACM SIGCOMM*, Aug., 2012.

[31] V. Giotsas, S. Zhou, M. Luckie, and k claffy. Inferring Multilateral Peering In *ACM CoNEXT*, Dec., 2013.

[32] A. Lodhi, N. Larson, A. Dhamdhere, and K. claffy. Using PeeringDB to Understand the Peering Ecosystem. *ACM SIGCOMM CCR*, 44(2), Apr., 2014.

[33] G. Gursun, N. Ruchansky, E. Terzi, and M. Crovella. Routing state distance: a path-based metric for network analysis. *ACM Internet Measurement Conference (IMC)*, Nov., 2012.

[34] M. Roughan, W. Willinger, O. Maennel, D. Perouli, and R. Bush. 10 Lessons from 10 Years of Measuring and Modeling the Internet's Autonomous Systems. *IEEE Journal on Selected Areas in Communications (JSAC)*, vol. 29, no. 9, 2011.

[35] P. Mahadevan, C. Hubble, D. Krioukov, B. Huffaker, and A. Vahdat. Orbis: rescaling degree correlations to generate annotated internet topologies. *ACM SIGCOMM*, Aug., 2007.

[36] A. Khan, T. Kwon, H.-c. Kim, and Y. Choi. AS-level topology collection through looking glass servers. In *ACM Internet Measurement Conference (IMC)*, Oct., 2013.

[37] Y. Zhang, R. Oliveira, Y. Wang, S. Su, B. Zhang, J. Bi, H. Zhang, and L. Zhang. A Framework to Quantify the Pitfalls of Using Traceroute in AS-Level Topology Measurement. *IEEE Journal on Selected Areas in Communications (JSAC)*, vol. 29, no. 9, Oct., 2011.

[38] E. Gregori, A. Improta, L. Lenzini, L. Rossi, and L. Sani. On the incompleteness of the AS-level graph: a novel methodology for BGP route collector placement. *ACM Internet Measurement Conference (IMC)*, Nov, 2012.

[39] Y. Shavitt, E. Shir. DIMES: Let the Internet measure itself. *ACM SIGCOMM CCR*, vol. 35, no. 5, Oct., 2005.

[40] H. V. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. Anderson, A. Krishnamurthy, and A. Venkataramani. iPlane: An Information Plane for Distributed Services. *USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, May, 2006.

[41] M. Lepinski and S. Kent. An Infrastructure to Support Secure Internet Routing. *RFC 6480*, Feb. 2012.

[42] traceroute.org, `http://www.traceroute.org`.

[43] traceroute.net.ru, `http://www.traceroute.net.ru`.

[44] bgp4.as, `http://www.bgp4.as/looking-glasses`.

[45] bgp4.net, wiki, `http://www.bgp4.net`.

[46] virus-net.ru, `http://www.virus-net.ru`.

[47] RIPE DB, ftp://ftp.ripe.net/ripe/dbase.

[48] RADB, ftp://ftp.radb.net/radb/dbase.

[49] RouteViews, http://www.routeviews.org.

[50] RIPE-RIS, http://www.ripe.net/ris.

[51] Packet Clearing House, https://www.pch.net/resources/data.php.

[52] Internet2, http://ndb7.net.internet2.edu/bgp.

[53] BGP potaroo, http://bgp.potaroo.net.

[54] Team Cymru WHOIS service, http://asn.cymru.com.

[55] BGP Toolkit, http://bgp.he.net.

[56] RIPE DB, ftp://ftp.ripe.net/ripe/dbase.

[57] RADb, ftp://ftp.radb.net/radb/dbase.

[58] AfriNIC FTP, ftp://ftp.afrinic.net/.

[59] APNIC FTP, ftp://ftp.apnic.net/.

[60] UCLA IRL, http://irl.cs.ucla.edu/topology.

[61] CAIDA ARK, http://www.caida.org/projects/ark.

[62] PeeringDB, http://www.peeringdb.com.

[63] Amsterdam Internet Exchange (AMS-IX), https://www.ams-ix.net/technical/specifications-descriptions/ams-ix-route-servers

[64] RouteViews, http://www.routeviews.org.

[65] RIPE-RIS, http://www.ripe.net/ris.

[66] Packet Clearing House (PCH), https://www.pch.net/resources/data.php.

[67] Internet2, http://ndb7.net.internet2.edu/bgp.

[68] NANOG, http://www.nanog.org.

[69] Tier-1 ASes, http://en.wikipedia.org/wiki/Tier1network.

[70] AfriNIC, http://www.afrinic.net.

[71] APNIC, http://www.apnic.net.

[72] ARIN, http://www.arin.net.

[73] LACNIC, http://www.lacnic.net.

[74] RIPE NCC, http://www.ripe.net.

[75] ISOLARIO, http://www.isolario.it.

[76] BGPmon, http://bgpmon.netsec.colostate.edu.

[77] EU Data protection Directive, http://europa.eu/rapid/press-release_MEMO-14-186_en.htm. Published: 12 Mar., 2014.

[78] R. GOVINDAN AND A. REDDY. An analysis of Internet inter-domain topology and route stability. In *IEEE INFOCOM*, 1997.

[79] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the Internet topology. In *ACM SIGCOMM*, 1999.

[80] K. Park and H. Lee. On the effectiveness of route-based packet filtering for distributed DoS attack prevention in power-law internets. In *ACM SIGCOMM*, 2001.

[81] G. Huston *et al.* Securing BGP: A Literature Survey, *IEEE Comm. Surveys Tutorials*, vol. PP, no. 99, 2010.

[82] R. Dingledine, N. Mathewson, and P. Syverson. Tor: the second-generation onion router. In *USENIX SECURITY*, 2004.

[83] X. Zhang, H.C. Hsiao, G. Hasker, H. Chan, A. Perrig, and D. G. Andersen. SCION: Scalability, control, and isolation on next-generation networks. In *IEEE Symposium on Security and Privacy*, May 2011.

[84] P. B. Godfrey, I. Ganichev, S. Shenker, and I. Stoica. Pathlet routing. In *ACM SIGCOMM*, August 2009.

[85] X. Yang, D. Clark, and A. W. Berger. NIRA: A new inter-domain routing architecture. *ACM/IEEE Trans. on Networking*, 15(4), Aug. 2007.

[86] J. Naous, M. Walfish, A. Nicolosi, D. Mazieres, M. Miller, and A. Sehra. Verifying and enforcing network paths with ICING. In *CoNext*, Dec. 2011.

[87] M. Lad, D. Massey, L. Zhang: Visualizing internet routing changes. *IEEE Transactions on Visualization and Computer Graphics*, Nov., 2006.

[88] A. Gupta, M. Calder, N. Feamster, M. Chetty, E. Calandro, E. Katz-Bassett *Passive and Active Measurement Conference (PAM)*, Mar. 2014.

[89] M. Lepinski *et al.* BGPSEC Protocol Specification, Network Working Group, Nov. 2013, Available: http://tools.ietf.org/html/draft-ietf-sidr-bgpsec-protocol-08

[90] D. McPherson, S. Amante, "Route Leak Attacks Against BGPSEC", Available: http://tools.ietf.org/pdf/draft-foo-sidr-simple-leak-attack-bgpsec-no-help-01

[91] C. Krugel, D. Mutz, W. K. Robertson, and F. Valeur, "Topology-Based Detection of Anomalous BGP Messages", in RAID, 2003, pp. 17–35

요약

인터넷의 Autonomous system(AS) 레벨 토폴로지는 미래의 프로토콜 설계, 성능 평가, 시뮬레이션 그리고 분석에 매우 중요하다. 과거 수십 년 동안 많은 연구가 진행되었지만, 인터넷의 AS레벨 토폴로지 전체를 완전히 파악하지 못하고 있다는 문제가 있으며, 심지어 최근 연구들은 이 문제들에 대해 예전보다 초점을 맞추고 있는 추세라고 한다. 따라서, 본 논문에서는 이러한 문제들의 해결을 위해 AS 레벨 토폴로지 데이터 소스들인 Looking Glass (LG) 서버와 Internet Routing Registries(IRR)에 초점을 맞추고자 한다.

LG서버는 110개국에 걸쳐 245개가 설치되어 있으며, 이 서버들에 쿼리를 보냄으로써 약 143K 개의 AS 링크들을 발견하여 AS 토폴로지를 추정하였다. LG 서버들로부터 AS 토폴로지에 있는 20K개의 새로운 AS 링크들을 발견하였으며, 이러한 링크들을 LG 서버들과 이웃한 620개의 AS들이 BGP Trace와 비교한 결과, 기존에 있는 RouteViews, RIPE-RIS, 그리고 PCH 에서 발견하지 못한 새로운 링크였다는 것을 발견하였다. 전체적으로, LG 서버들로부터 BGP Trace 들을 모으는 것은 기존 BGP Collectors의 한계를 보완할 수 있지만, LG 서버에서 발견되는 AS 토폴로지는 제한된 Vantage points 와 이웃한 AS에 의해 사용되고 있는 BGP export 정책들 때문에 어려운 점이 있다.

장점들도 있지만, LG 서버들의 이러한 한계점을 보완하기 위해 Internet Routing Registries (IRR)을 이용하였다. IRR은 도메인 간 라우팅 정책들을 등록시키기 위하여 AS에 사용되는 데이터베이스이다. 본 논문에서는 IRR로부터 AS 레벨 토폴로지를 추출하는 방법을 처음으로 소개하였으며, bilateral 그리고 multilateral peering 링크를 이용하였다. 2013년 12월 1일부터 모은 IRR 데이터를 바탕으로, 610K개의 AS 링크들을 추출하였으며, 이 중 68%는 BGP, traceroute, Internet eXchange points (IXPs) cliques에서 이미 발견된 것들이다. IXP의 멤버 AS들을 이용하여 크고 작은 여러 IXP들의 Peering 행렬들을 파악하는 것이 가능하다. 따라서, 본 논문에서는 IRR에 저장되어 있는 라우팅

정책을 이용하는 AS들 간의 비즈니스 관계를 추론하는 방법을 본 논문에서 소개하고자 한다. 이 알고리즘의 정확도는 현재 존재하는 BGP의 AS PATH를 이용하여 AS 관계들을 유추하는 알고리즘들과 비교해도 경쟁력이 있다는 것을 보여준다. IRR을 이용하는 것은 인터넷에서 상당히 많은 AS들에 의해 사용되기 때문에, 이것을 이용하는 것이 인터넷의 구조, 성능, 유동성, 진화들에 대해 더 높은 이해도를 위한 상호보완적인 방법이라고 결론 내릴 수 있다.

주요어: 인터-도메인 라우팅, 루킹 글래스 서버, 인터넷 라우팅 레지스트리

학번: 2008-23532