Public Review for

A Comparative Study on IP Prefixes and their Origin ASes in BGP and the IRR

Akmal Khan, Hyun-chul Kim, Ted “Taekyoung” Kwon,
and Yanghee Choi

This paper undertakes the quantification of the validity of IRR data along various dimensions: overall consistency with BGP data, dependence on type of AS, and on IRR region which is checked. Specific focus is on the (IP-prefix, AS origin) pair, PO pair for short, the consistency of which is checked across the IRR data and the BGP data. For the IRR data, a representative subset of IRR data is taken from two IRR databases, RIPE and RADB, in Jan 2013. And for the BGP data, PO pairs published by CAIDA (in the same month), as extracted from RouteViews, is taken. Using this data, the paper shows that registration of PO pairs in IRR is common among more than three-quarters of ASes.

Such registration is especially common among small-to-medium sized transit providers, likely since they manage a small number of customer ASes (a few tens to a few hundreds). The analysis also shows that the trustworthiness of the IRR data varies with region; some regional registries are better maintained than others. The significance of this work lies in the potential application of IRR data in preventing or at least mitigating inter-domain routing issues due to prefix misconfigurations or prefix hijacking attacks.

While the paper is an initial step in quantifying the validity of IRR data, the actual scheme for the use of IRR for BGP prefix filtering is part of this paper’s future work. Also of potential interest in related future work, is the examination of inter-AS relations in the context of overlapping PO pairs in the IRR, as well as checking the overall conclusions in the paper using other IRR data and using active measurements (e.g. using PlanetLab).

Public review written by
Bhaskaran Raman
Department of CSE, IIT, Bombay, India

A Comparative Study on IP Prefixes and their Origin ASes in BGP and the IRR

Akmal Khan, *Hyun-chul Kim, Ted "Taekyoung" Kwon, and Yanghee Choi
Seoul National University,*Sangmyung University
akmalshabaz,hyunchulk@gmail.com, tkkwon,yhchoi@snu.ac.kr

ABSTRACT

The IRR is a set of globally distributed databases with which ASes can register their routing and address-related information. It is often believed that the quality of the IRR data is not reliable since there are few economic incentives for ASes to register and update their routing information in a timely manner. To validate these negative beliefs, we carry out a comprehensive analysis of (IP prefix, its origin AS) pairs in BGP against the corresponding information registered with the IRR, and vice versa. Considering the IRR and BGP practices, we propose a methodology to match the (IP prefix, origin AS) pairs between those two datasets. We observe that the practice of registering IP prefixes and origin ASes with the IRR is prevalent, though the quality of the IRR data varies substantially depending on routing registries, regional Internet registries (to which ASes belong), and AS types. The IRR can help improve the security level of BGP routing by making BGP routers selectively rely on the corresponding IRR data considering these observations.

Categories and Subject Descriptors

C.2.2 [**Network Protocols: Routing Protocols**]:

Keywords

Inter-domain Routing, BGP, IRR

1. INTRODUCTION

The Internet Routing Registry (IRR) is a set of databases that are to be used by Autonomous Systems (ASes) to register their inter-domain routing information. Although the idea of the IRR usage has been around since the NSFNET era [1], the IRR has often been considered outdated since many ASes consider their routing policies and peering information private, and hence they may not publicize such information [3, 5]. In line with other studies [3, 10], our initial investigation confirms that the detailed routing policies such as a number of customer and provider ASes of an AS are not registered mostly. However, (IP prefix, origin AS) pairs of a vast majority of ASes are registered with the IRR. Figure 1 shows the number of (IP prefix, origin AS) pairs (or shortly PO pairs) in the IRR as well as in BGP from Nov. 2010 to Feb. 2013.

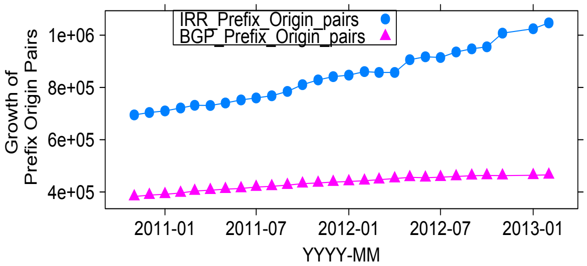

Figure 1: The number of (IP prefix, origin AS) pairs in the IRR and BGP (from RouteViews collectors) are plotted respectively from Nov. 2010 to Feb. 2013.

While different growth patterns of the PO pairs between BGP and the IRR are evident, there has been no study that explains (i) how consistent are the PO pairs of the IRR compared with those of BGP (and vice versa)? (ii) which (sort of) ASes register their PO pairs with the IRR or not, (iii) how many transit ASes enforce the IRR registration on their customer ASes[1], and (iv) how the practice of the IRR registration of PO pairs varies across different types of ASes and regional Internet registries (RIRs).

Finding empirically-grounded answers to these questions is important as it helps the research and operational community whether and how much the IRR can help mitigate or prevent wrong IP prefix announcements over BGP. That is, if ASes can trust the information stored in the IRR, then inter-AS filters (of IP prefixes) can be created to prevent or mitigate prefix misconfigurations and prefix hijacking attacks [3, 8, 9, 12].

To understand the registration practice of IRR PO pairs depending on (i) AS types, (ii) RIRs (to which ASes belong), (iii) origin/transit ASes, and (iv) routing registries (RRs), we conduct a comprehensive study on the comparison of the PO pairs in BGP against the IRR ones. As similar to other studies [8, 9], we find matching PO pairs between the IRR and BGP, which means that we check whether the same PO pair can be observed both in the IRR and BGP. However, for the same IP prefix, its origin AS found in BGP and the one registered with the IRR can be different. This is partially due to the practice of transit ASes registering the PO pairs of their customer ASes with the IRR on behalf of

Corresponding Authors: Ted “Taekyoung” Kwon (tkkwon @snu.ac.kr) and Hyun-chul Kim (hyunchulk@gmail.com)

[1] There has been only anecdotal evidence that many ASes require their peer or customer ASes to register their PO pairs [4, 6].

their customer ASes. Thus, we propose a methodology to check such cases from the relevant information in the IRR. Moreover, the origin AS mismatches between BGP and the IRR can also be due to the inter-domain routing practice in BGP such as prefix aggregation or deaggregation.

The key observations from the comparison of the PO pairs between BGP and the IRR in Jan. 2013 are as follows.

- Around 82% of (475 K) IPv4 BGP PO pairs were found in the IRR. For around 4% of the BGP PO pairs, the same IP prefixes were found in the IRR but they were with different origin ASes, due to misconfigured or malicious IP prefix announcements or possibly stale information in the IRR. For the remaining 14% of the BGP pairs, we could not find a matching (either the same, sub, or super) IP prefix from the IRR ones at all (§ 4).
- We observe that the practice of registering PO pairs with the IRR is prevalent; 78% of ASes register all of their BGP-announced IP prefixes with the IRR; around 11% of ASes register a subset of their BGP prefixes with the IRR. The remaining 11% ASes register none of their BGP prefixes with the IRR; they are mostly single-homed Stub ASes of a few large transit providers. Overall, ASes belonging to RIPE, APNIC, and AfriNIC register their prefixes with the IRR more actively than ASes in ARIN and LACNIC (§ 4).
- Our analysis on the IRR registration practice depending on AS types reveals that small transit providers have more actively adopted the IRR registration than large transit providers, most likely due to the smaller number of their customer ASes (§ 4).
- We investigate individual RRs and find that the trustworthiness of the PO pairs varies substantially across them. The PO pairs in well-maintained RRs can be used for the purpose of BGP message verification. We find that JPIRR (the NIR of Japan) and BELL (BELL Canada) are the best maintained ones. Considering the large numbers of PO pairs, RADB, RIPE, and NTTCOM are also relatively well-maintained large RRs (§ 5).
- Based on our findings, we conclude that (i) the IRR can help mitigate the prefix origin based security vulnerability [3] of BGP, and (ii) the IRR is readily available and possibly complementary to Resource Public Key Infrastructure (RPKI) [7][2].

The rest of the paper is organized as follows. In Section 2, we present the background on the IRR. Section 3 describes our methodology and datasets. Section 4 shows the analysis of BGP pairs with reference to the IRR ones. Section 5 presents the analysis of IRR pairs against BGP ones. Section 6 presents our validation efforts. Sections 7 and 8 discuss implications and related work, respectively. Finally, we make concluding remarks in Section 9.

2. IRR

As of Feb. 1st, 2013, the IRR consists of 34 routing registries (RRs) mostly maintained by either RIRs[3], national Internet registries (NIRs), local Internet registries (LIRs), or Internet service providers (ISPs). An RR refers to a database that stores routing policy information of ASes such as IP prefixes originated by ASes and routing policies towards their neighbor ASes. The routing policy information in the IRR is expressed by a standard language, Routing Policy Specification Language (RPSL) [2]. The RPSL defines several kinds of objects, most of which can be classified into three groups: (i) *inetnum or inet6num* objects describe IPv4 or IPv6 address allocation, (ii) *route, route6, aut-num, route-set, as-set* objects describe the routing policies, and (iii) *mntner, person, and role* objects describe who administers the routing policies and so on.

Table 1: **route** and **mntner** objects in the IRR; an example.

mntner: MNT-AS1	**route:** 10.1.1.0/16
mnt-by: MNT-AS1	**origin:** AS1
changed: 20130101	mnt-by: MNT-AS1
source: RR2	changed: 20130101

An RPSL object is represented as a list of attribute-value pairs as illustrated in Table 1, where key attributes are written in boldface. For example, **mntner** is a key attribute of a maintainer object, **route**, and **origin** are key attributes of a route object . Some common attributes across objects are: the *mnt-by* attribute specifies the maintainer of a given object; the *changed* attribute provides the last-updated date; and the *source* attribute specifies the name of the RR where the RPSL object is registered. An AS starts registering its routing policy with an RR by first making a request to create a maintainer account to the RR. The maintainer information of an AS is stored in a *mntner* object, which is the authorized entity to add, delete, or modify objects. Once a *mntner* object is created, the maintainer of an AS is allowed to register necessary RPSL objects with the RR.

3. METHODOLOGY AND DATASETS

We present our methodology to compare PO pairs between the IRR and BGP.

3.1 IRR and BGP PO pairs

The IRR data was downloaded from the RIPE RR [14] and RADB [15] in Jan. 2013. We extract 1.24 M IPv4 PO pairs from the route objects (ROs). We then removed 73 PO pairs registered by private ASNs, 250 for private IP addresses[4], and 210 K duplicate PO pairs[5] from the IRR data. Finally, around 1.03 M PO pairs belonging to 43.5 K ASes are selected from the IRR data, which are collectively referred to as *IRR pairs*.

We downloaded the BGP PO pairs [16] published by CAIDA in Jan. 2013, which are extracted from the routing tables and update archives of RouteViews [17] BGP collectors. We

[2] RPKI [7] has just started its deployment, and is still in its initial deployment phase.

[3] The five operational RIRs are AfriNIC [28], APNIC [29], ARIN [30], LACNIC [31], and RIPE NCC [32]

[4] The reason why private IP prefixes (and private ASes) are registered in the IRR is that the IRR is used by ASes to manage their customer ASes. So it is for operational reasons that ASes register private IP prefixes or ASes in the IRR. For example, a private IP prefix/AS is used by a customer AS with the approval of its provider AS and that private IP prefix/AS is only used to communicate between the provider and customer ASes.

[5] An AS may register its PO pairs with multiple RRs if it has multiple provider ASes that use different RRs respectively.

found around 475 K pairs that belong to 43.6 K ASes. These BGP PO pairs are collectively referred to as *BGP pairs*[6].

3.2 Comparison of BGP and IRR Pairs

In this section, we present our proposed methodology for checking matching PO pairs between the IRR and BGP. Note that the same methodology will be applied once from BGP perspective in Section 4 and then again from the IRR perspective in Section 5. For the sake of brevity, we describe the methodology from BGP perspective only.

1. Direct Prefix Origin Match (POM). If the origin AS in a PO pair in BGP matches with the one in the IRR pair directly, it is called direct Prefix Origin Match (Direct-POM). Even if the prefix lengths of the two PO pairs (i.e., BGP and IRR pairs) are different (i.e., less specific or more specific prefix match), the PO pairs are deemed matching if their origin ASes are the same since an AS can announce different lengths for a given IP prefix (e.g., due to traffic engineering practice [11]).

2. Mntner Prefix Origin Match (POM). Recall that a maintainer can maintain the ROs of multiple ASes in the IRR. Therefore, we examine *mntner* objects and related RPSL objects whose *mnt-by* attributes have the same maintainer information. For example, consider a case where BGP shows a prefix 10.0.0.0/19 originated by AS2, while the IRR has the same prefix 10.0.0.0/19 registered by AS3. If we observe in the IRR that the PO pairs of AS2 and AS3 are maintained by the same maintainer, those two pairs are deemed matching (Mntner POM). This is typically called *Proxy Registration*, which is the practice of provider ASes registering PO pairs of their customer ASes with the IRR, which can be due to service level agreements (SLAs) between provider and customer ASes.

Mntner data. Though there are 43 K maintainers in the IRR, only 25 K maintainers are found to register their PO pairs with the IRR. We observe that 21 K maintainers maintain ROs of a single AS, while the remaining 4 K maintainers register ROs of multiple ASes (ranging from two to 4 K ASes). We also found that dozens of maintainers register their data across different RRs. For instance, *LEVEL3-MNT* maintains PO pairs of around 4 K ASes across LEVEL3 and RIPE RRs.

3. AS Link Prefix Origin Match (POM). Even when Direct POM and Mntner POM methods fail, the difference between the origin AS in BGP and the mismatching one in the IRR can be due to legitimate reasons: (i) prefix aggregation [13], (ii) prefix deaggregation [11], and (iii) static routing [11]. These three cases (i.e., inter-domain routing practice by ASes) can be checked by leveraging inter-AS link information. For example, consider a case where AS1 registers 10.0.0.0/24 prefix with the IRR, while 10.0.0.0/19 is advertised by AS2 in BGP. If an AS-level link between AS1 and AS2 is found in any inter-AS link data extracted from BGP and traceroute measurements, then we conclude that AS1 is a legitimate origin AS of the corresponding prefix (AS Link POM).

AS link data. There have been numerous efforts to construct the AS-level topology (i.e., a collection of AS-level links) based on BGP trace, traceroute, and the IRR data. Unfortunately, the completeness of the inferred AS topology remains elusive since the Internet is operating in a decentralized manner and continuously changing [5]. To use comprehensive AS links[7] that are available in public, we build a dataset of AS links by combining the following three data sources: (i) the UCLA IRL topology data [19] which is based on BGP traces, (ii) traceroute-based measurements published by DIMES [20], CAIDA Ark [21], and iPlane [22], and (iii) AS links derived from the IRR *aut-num* objects. All of the AS link data have been collected for 31 days, from Jan. 1st to 31st, 2013, except DIMES, whose most recent publicly available dataset was collected and downloaded in Apr. 2012.

While processing the above data, we found overlapping AS links across the data sources. Table 2[8] indicates the level of overlapping between the inter-AS link data. Nevertheless, each AS link dataset contributes unique AS links that are combined to build a more complete view of the AS-level topology of the Internet, which consists of around 445 K AS links in total.

Note that we have decided to use the recently (and regularly) published AS topology datasets only, not including ones such as chen *et al.*'s [13], which had been collected using BitTorrent P2P clients in 2007-2008. Since we can not quantify how much of this AS topology dataset is outdated, we decided not to use it.

Table 2: The number of overlapping and unique (in bold) inter-AS links between various AS link datasets. All of the AS link data have been collected for 31 days, from Jan. 1st to 31st, 2013, except DIMES, whose most recent publicly available dataset was collected and downloaded in Apr. 2012.

Dataset	#links	IRR	Ark	iPlane	DIMES	UCLA
IRR	203 K	**129 K**	32 K	20 K	27 K	67 K
Ark	141 K	32 K	**57 K**	47 K	53 K	58 K
iPlane	76 K	20 K	47 K	**17 K**	44 K	38 K
DIMES	105 K	27 K	53 K	44 K	**34 K**	43 K
UCLA	182 K	67 K	58 K	38 K	43 K	**78 K**

3.3 AS Types and RIR data

AS Types. We classify ASes into four categories: Tier-1, Large Transit Provider (LTP), Small Transit Provider (STP), and Stub ASes. We find the list of 16 Tier-1 ASes from [27]. On the basis of the number of PO pairs announced in BGP by a given AS, the given AS is classified into one of the three categories: a Stub AS if it has ≤ 5 PO pairs, Small Transit Provider (STP) if it has > 5 and $\leq$ 1 K PO pairs, Large Transit Provider (LTP) if it has > 1 K PO announcements. We use these thresholds based on our analysis results on the number of IP prefixes announced by ASes in BGP. Consequently, we classify the total of 43.6 K

[6] We also experimented with RIPE-RIS [18] BGP traces but they do not add that much new information to what has already been observed in the RouteViews BGP traces.

[7] AS links refer to AS-level links.

[8] Note that we choose not to use the AS relationship information as currently only 24% of our combined AS-level link data has such information. The UCLA dataset contains AS relationship information as a part of their AS-level link data set, none of the other datasets provide such information. Thus, incorporating the AS relationship info into our methodology and quantifying its impact on the results are left as future work.

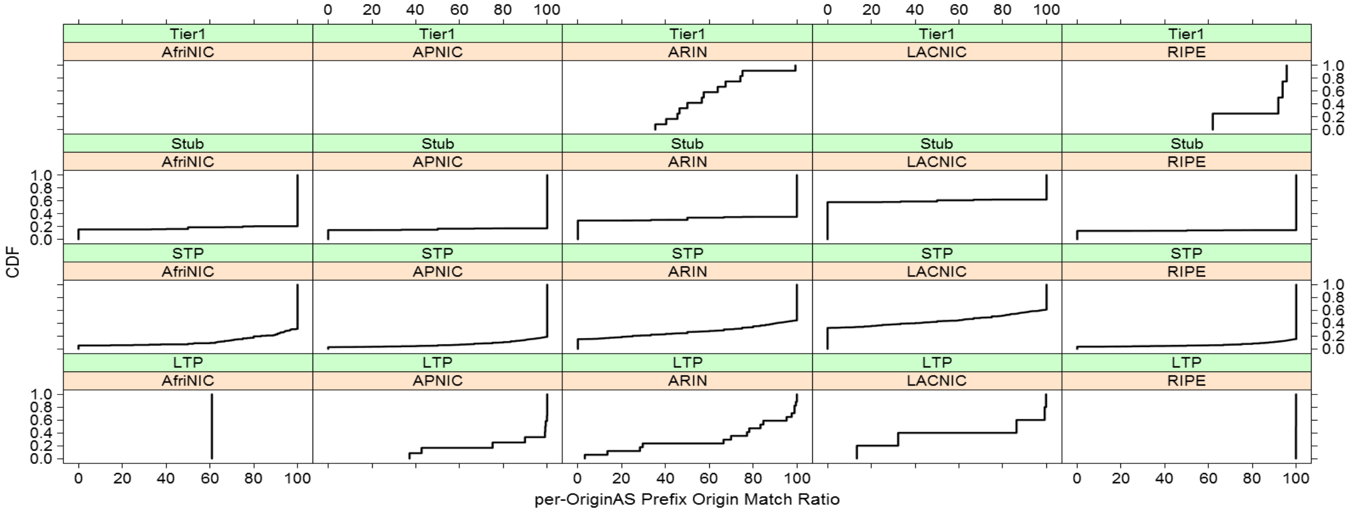

Figure 2: CDFs: per Origin-AS POM ratios across AS types and RIRs.

ASes found in the BGP PO pairs into 16 Tier-1, 37 LTP, 10.7 K STP, and 32.7 K Stub ASes.

RIR's IP allocations. To map IP prefixes to the corresponding RIR regions, we downloaded the IP allocation records from the five RIRs in Jan. 2013.

4. ANALYSIS OF BGP PO PAIRS

We apply the methodology described in Section 3.2 on the BGP PO pairs (475 K) against the IRR pairs (1.03 M). Then we analyze the prefix origin match (POM) results from various perspectives (e.g., AS types and RIR regions).

How many (IP prefix, origin AS) pairs in BGP are found in the IRR PO pairs? Around 82% BGP PO pairs are found in the IRR ones, with 71% of Direct POM, 4% of Mntner POM, and 7% of AS Link POM. The largest contribution of Direct POM indicates that a large percentage of BGP pairs are directly consistent with the IRR ones, while Mntner POM and AS Link POM also exhibit non-negligible percentages, which are detailed as follows.

Mntner POM. We observe that around 19 K BGP pairs (4%) match with the IRR ones by Mntner POM. There are 4 K (i.e., 16 % of all the 25 K maintainers who register their PO pairs with the IRR) maintainers who register their PO pairs on behalf of their customer ASes. For example, large provider ASes like Level-3 (AS 3356) and PCCW Global (AS 3491) typically register around 1 K pairs on behalf of their customer ASes.

AS Link POM. A total of 9 K AS links are used to find 33 K (7%) BGP pairs matching with IRR ones. We find that around 7 K pairs match using 3 K AS links, which are found in only one of the AS link data such as DIMES (1 K), Ark (1 K), and IRR (0.4 K). The remaining 26 K BGP pairs match with the IRR ones by using the other 6 K AS links that are found in more than one AS link datasets. For example, 920 AS links are common between DIMES and Ark. 346 links are common among DIMES, Ark, and iPlane. It is interesting to note that 3.5 K (out of 9 K in total) AS links are observed only from the traceroute-based AS link datasets. While the accuracy of the traceroute-based AS level topology is often questioned (e.g., [23]), the above result suggests that the AS Link POM method can also be used to validate the AS links discovered using traceroute-based measurements.

Prefix origin mismatches between BGP and IRR. We found around 4% (19 K) of BGP pairs that have the same prefixes in the IRR but with mismatching origin ASes. The largest contribution of origin AS mismatches (12 K out of 19 K) come from /24 mask length prefixes, which seem to have been outdated as these PO pairs have not been updated since 2005-2009. Our analysis shows that /24 blocks are more often used by small ASes, which due to mergers or acquisitions of these ASes are returned to the IRR and then allocated to other ASes.

Yet, outdated IRR records may not be the only reason for origin AS mismatches, since we also observed 1 K (out of 19 K) mismatching BGP pairs whose corresponding IRR records have been updated very recently, in between Jun. 2012 and Jan. 2013. We believe that future research in the area of building more complete AS topologies and understanding BGP routing policies will shed more light on possible reasons for such origin AS mismatches between the IRR and BGP.

Per Origin-AS IRR registration. We further analyze POM results depending on the origin ASes announcing those prefixes in BGP. Such analysis can help us look into the prefix registration practice of ASes. Per Origin-AS POM is defined as the ratio of the number of PO pairs (in BGP) of an AS matching with the IRR ones to the total number of PO pairs announced by the AS. Figure 2 shows the CDF charts of per Origin-AS POM results, which varies depending on AS types and RIRs: (i) Stub ASes show higher per Origin-AS POM as compared to STPs and LTPs since Stub ASes have to manage far fewer PO pairs than STPs and LTPs, and (ii) RIPE, APNIC, and AfriNIC exhibit high per Origin-AS POM indicating the practice of good IRR registration while ARIN and LACNIC show low per Origin-AS POM.

Overall, 34 K (78% out of 43.6 K) ASes appearing in BGP across RIRs exhibit 100% PO pairs registration in the IRR. Furthermore, we find 3.4 K (8% of 43.6K) ASes do not register any PO pairs with the IRR. We observe that most of

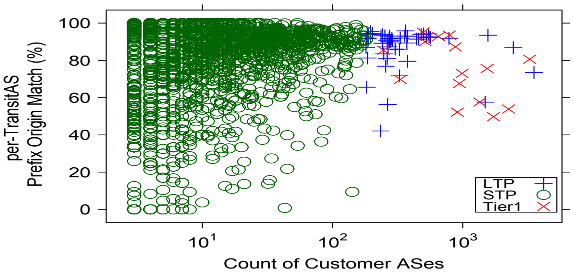

Figure 3: Per Transit-AS POM results, calculated by averaging those of their customer ASes.

these ASes are single-homed customers of Tier-1 and LTP ASes. That is, some Tier-1 and LTP ASes do not seem to require their customer ASes to make the IRR registration. Overall, we observe that ARIN and LACNIC have a high percentage of ASes with no IRR registration. While the detailed results are omitted, we highlight that there are 1.3 K (3% of around 43.6 K) ASes whose PO pairs are announced in BGP but none of their prefixes have matching Origin ASes in the IRR, possibly due to various operational reasons such as IP address space transfers across ASes and merger, acquisition, or closure of an AS, etc.

We observe that 4.8 K (11 % out of 43.6 K) ASes make partial IRR registration, i.e., only a proper subset of their BGP PO pairs have been found in the IRR. It can happen as an AS may have multiple provider ASes, each of which may have different requirements on the IRR registration. Note that all the Tier-1 ISPs register only a subset of their IP prefixes with the IRR although the Tier-1 ISPs typically originate a very small number of prefixes in BGP (since their main role is transiting, not hosting).

Per Transit-AS IRR registration. From the analysis of the per Origin-AS POM results of BGP pairs, we find that a large number of ASes register their PO pairs with the IRR. However, it is not clear whether the registration is a voluntary activity or influenced by their transit ASes. We analyze the POM results as per transit ASes; that is, ASes carrying the prefixes of their customer ASes towards the core of Internet. To calculate the POM ratio of transit ASes, we extract around 6.5 K transit ASes from the UCLA IRL data [19]: STP (6.4 K), LTP (48), and Tier-1 (16). Then, the POM ratio of a transit AS is calculated by taking the average of the POM results of its customer ASes.

Figure 3 shows the POM ratio of the transit ASes. We observe that most of the STPs have high POM results especially as the number of their customer ASes increases. This suggests that these STPs find the IRR useful in managing PO pairs of their customer ASes. However, a slight decrease in the POM values of LTPs is noticeable with the increase in number of their customer ASes, which suggests that transit ASes with thousands of customers may face more difficulties in maintaining their IRR entries. Overall, the IRR is more popular in the realm of small (or mid-sized) transit providers, since they manage a smaller number of customer ASes; typically tens to a few hundreds number of customer ASes.

5. ANALYSIS OF IRR PO PAIRS

We apply the same methodology as described in Section 3.2 on the IRR pairs (1.03 M) against the BGP pairs (475 K). Such analysis is important to answer the following research questions:

How many (IP prefix, origin AS) pairs in the IRR are found in BGP PO pairs? Around 87% IRR PO pairs match with the BGP ones: Direct POM (65%), Mntner POM (12%), and AS Link POM (10%). We find 5% of the PO pairs whose prefixes match with the ones in BGP, but their origin ASes mismatch. The remaining 8% of the PO pairs are not found in BGP at all. We note that the different growth patterns of BGP and IRR pairs observed in Figure 1 are due to the prevalent registration of more specific PO pairs in the IRR. According to our analysis, as much as 50% of IRR pairs match with the BGP PO pairs that have the less specific prefixes.

To investigate the reasons of the 8% of the IRR PO pairs that are not observed in BGP, we downloaded PO pairs from the CAIDA BGP data [16], starting from Jan. 2009 to Dec. 2012 (BGP history time window), and checked whether those IRR pairs have ever been advertised in BGP. We found that around 55% of the unobserved prefixes are historical, that is, they used to be advertised in BGP in the past. The remaining 45% may have been advertised before Jan. 2009 or yet to be announced.

To further investigate, we sent an IRR usage survey questions to network operators on the NANOG (North American Network Operators' Group [25]) mailing list. We also emailed our IRR analysis results (along with the IRR usage survey questions) to 200 network operators that are found to have differences in what these ASes register in the IRR and what is visible of these ASes in BGP, only 50 out of which have replied to us. While these operators have shared with us the reasons of using the IRR and their usage practices, they do not want to comment on their routing practices related to customer ASes. Due to no input from the operators on validation aspects of our work, we discuss only the two findings from from our IRR usage survey as follows:

Why ASes keep historical pairs? The two main reasons for historical PO pairs are as follows: (i) Once registered, many ASes do not care much about removing the data from the routing registries (RRs). For instance, when ASes move from one RR to another due to the change of their provider AS, as different providers of an AS can require the AS to use different RRs. (ii) Many ISPs manage their customer's PO pairs in the IRR on behalf of them, and they often just keep the PO pairs of their customers even if those customers are no longer using their services. Just in case the customer wishes to resume the services from the ISP, then the ISP would not have to re-enter the information in the IRR, which helps in providing quick transit services to the returning customers.

Why ASes register unannounced pairs? There is no time restriction between the date of IRR registration and the date of its first announcement in BGP; the latter can happen even after years. For instance, ASes can reserve the allocated IP prefixes for future use. IRR pairs can also be kept being unannounced in BGP if ASes are using them for internal network management.

Per RR Analysis. We have also investigated how POM ratios vary across different RRs. Figure 4 shows POM ratios of the 14 large RRs (each with more than 1 K PO pairs) in

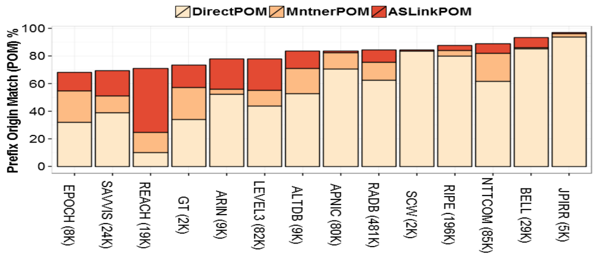

Figure 4: POM ratios across the 14 large RRs that are with more than 1 K PO pairs registered, along with the number of their PO pairs.

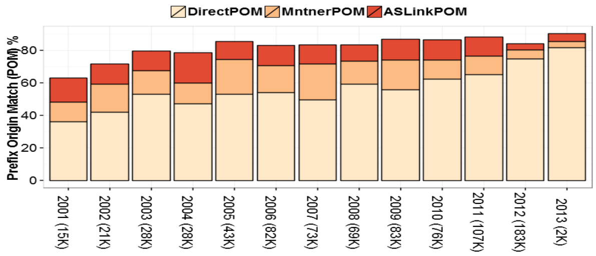

Figure 5: POM ratios according to the last-updated date.

increasing order of the POM ratio values. The x-axis shows the name of the RR along with its count of PO pairs. The y-axis indicates what percentage of PO pairs of an RR are also found in BGP. We find that JPIRR (the NIR of Japan) and BELL (Bell Canada) exhibit high portion of Direct POM. In many RRs, a substantial portion of POM is due to Mntner POM and AS Link POM, which highlights the effectiveness of the two proposed POM methods.

POM ratios per (last-updated) year. We also look into whether the last update dates of the IRR PO pairs have any relation with their POM results. In particular we seek to validate the common belief that the older IRR pairs are more likely to have been outdated. For instance, Oliveria et al. [5] discarded IRR data older than one year setting up their experimental environment, with the assumption that IRR data not updated for more than a year is highly likely to be outdated. Figure 5 shows that as the last update year of a PO pair becomes older, its origin AS shows an increasing tendency of mismatching with the one in the BGP trace. However, we also observe that IRR PO pairs which have not been updated for several years (e.g., 63 % in the year 2001) are still found in BGP, which suggests that IRR records should not be assumed to have been outdated just because they are several years old. The steady portion of Mntner POM over the years indicates that the IRR PO pair registration by a provider AS on behalf of its customer ASes has been practiced for many years. Finally, the contribution of AS link POM over the years also highlights its effectiveness.

6. VALIDATION

The most challenging part of any Internet routing-related study is to validate results in the absence of the ground truth [5, 13, 23]. Despite the challenge, we have performed validation of our proposed IRR origin AS checking methodology, by using the ARIN origin AS dataset (AOAS) [26]. We downloaded the AOAS data in Jan. 2013, which has around 367 K PO pairs.

Results. Around 41 K IRR PO pairs are validated by AOAS, that is, the origin ASes of these IRR pairs are the same as the corresponding ones in AOAS. Note that their origin ASes in BGP pairs can be different. Most of the validated IRR pairs are found by Direct POM (36 K pairs in 4 K ASes). Around 2 K and 3 K IRR PO pairs are validated by Mntner POM and AS Link POM, respectively. Also, those IRR pairs verified by each method belong to 200 ASes and 450 ASes, respectively, which signifies that each origin AS checking method is quite effective in finding matching pairs. Interestingly, we also find that 2 K IRR pairs that were not found in BGP PO pairs (in Section 5) by our methodology have matching ones in AOAS.

Given the lack of comprehensive high quality topological and inter-domain routing data, our validation is admittedly of limited quality and coverage. Yet, this section demonstrates that our proposed methodologies are effectively working, at least according to the validation results obtained with a small ground truth dataset, which is currently provided by only ARIN. We anticipate that other RIRs will follow ARIN in future and start publishing the ground truth datasets, which are submitted to RIRs by ASes when applying for the IP/AS number resources. Moreover, we are working on improving the validation aspects of our work using Looking glass (LG) servers. More specifically, if a (prefix, origin AS) pair is in the IRR but not found in BGP (due to prefix aggregation), a Looking glass server can be used as well to validate the IRR PO pair, as BGP traces from RouteViews and RIPE-RIS may only provide aggregated prefixes, whose origin AS is different from the IRR one.

7. RELATED WORK

Neighborhood Watch [8] and TERRAIN [9] have been proposed to use the IRR to improve BGP security against prefix origin attacks. TERRAIN and Neighborhood Watch focus on the consistency across different kinds of RPSL objects maintained in the IRR. That is, both approaches require the existence of multiple routing policy objects (*aut-num, organization, inetnum*, etc.) as well as the consistency among those objects to decide whether an incoming BGP UPDATE message has a valid PO pair. However, those RPSL objects are not usually maintained in the RRs; most ASes are interested in registering only *aut-num* and *route* objects [24]. Thus, approaches like TERRAIN and Neighborhood can be applied only to RRs with the detailed routing policy information, such as RIPE. On the other hand, our work requires only PO pair data, which reflects the common RR registration practice of ASes across all the RRs. Moreover, we propose methods to check different origin ASes between BGP and the IRR considering the practice of BGP announcements such as route aggregation.

Siganos *et al.* [10] developed Nemecis to analyze routing policies of ASes by extracting routing policy information from the *import* and *export* attributes of their *aut-num* objects in the IRR. By extracting neighbor information from the *aut-num* objects, Nemecis generates AS relationship data, which is then checked whether they are consistent with those from BGP traces. On the other hand, our work primarily focuses on extracting PO pairs from *route* objects in the IRR, as there is anecdotal evidences that provider ASes often ask their customer ASes to register *route* objects with the IRR, while registration of *aut-num* objects are not usually asked [4, 6]. Also, even the minimal routing policy information like PO pairs can help mitigate BGP security problems such as misconfigured or spoofed BGP announcements [3].

8. DISCUSSION AND FUTURE WORK

Our comparative study highlights that in contrast to common negative claims, registration of BGP PO pairs in the IRR is prevalent, which gives us opportunities to build important services based on the IRR data.

IRR for BGP prefix filtering service. IRR data can be used to validate BGP Prefix-Origin AS pairs. A similar proposal known as RPKI is underway to ensure the integrity of BGP messages for secure inter-domain routing [7]. Considering that RPKI would take many years to be fully functional, we stress the role of the IRR as an important viable alternative source of trustworthy PO pairs, which can help thwart IP prefix origin AS based hijacking attacks and BGP misconfigurations.

As the IRR does not cover the whole IP space and its largest missing part belongs to ARIN (§ 4), we propose to augment the IRR by the ARIN origin AS dataset [26]. In addition, as POM ratios substantially vary across different RRs, the IRR can help improve the security level of BGP routing by making BGP routers selectively rely on the corresponding IRR data, considering these observations.

IRR as another source of AS-link data. IRR's role in building more complete AS-level topology is highlighted in table 2. That is, 129 K unique AS links exist only in the IRR even when compared against the comprehensive combination of all the other publicly available datasets used in this paper [19–22]. Further analysis and validation of our combined AS level topology (of 445 K AS links in total in § 3) is needed. For example, to find out how such a comprehensive AS-level topology helps in detecting AS path spoofing attacks in BGP [3].

9. CONCLUSION

To investigate the registration practice of (IP prefix, origin AS) pairs in the IRR, we first comprehensively consolidated Internet routing and address-related data: BGP traces, AS-level links from BGP and traceroute-based measurements, and so on. The consolidated data is then used to analyze the (prefix, origin AS) pairs registered with the IRR against the corresponding ones announced in BGP, and vice versa. We find that the practice of registering (prefix, origin AS) pairs in the IRR is prevalent. However, the quality of the IRR data can vary significantly depending on RRs, RIRs (to which ASes belong), and AS types. We argue that, as similar to other studies [8, 9], the IRR can help improve the security level of BGP routing by making BGP routers selectively rely on the corresponding IRR data considering these factors. As a part of our on-going work, we are designing and implementing a BGP prefix filtering service based on the IRR.

10. ACKNOWLEDGEMENTS

This research was supported in part by the KCC (Korea Communications Commission), Korea, under the R&D program supervised by the KCA (Korea Communications Agency) (KCA-2013-11-911-05-002), in part by the Basic Science Research Program through the National Research Foundation of Korea (NRF), funded by the Ministry of Education (No. 2013-R1A1A2010474).

11. REFERENCES

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

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

[3] 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.

[4] Route Filter Panel. *Nanog43*, Jun. 2008. http://www.nanog.org/meetings/nanog43.

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

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

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

[8] G. Siganos and M. Faloutsos. Neighborhood Watch for Internet Routing: Can we improve the Robustness of Internet Routing Today? In *IEEE INFOCOM*, May. 2007.

[9] K. Sriram, O. Borchert, O. Kim, P. Gleichmann, and D. Montgomery. A Comparative Analysis of BGP Anomaly Detection and Robustness Algorithms. In *IEEE CATCH*, Mar. 2009.

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

[11] L. Cittadini, W. Muhlbauer, S. Uhlig, R. Bush, P. Francois, and O. Maennel. Evolution of Internet Address Space Deaggregation: Myths and Reality. *IEEE JSAC*, vol. 28, no. 8, Oct. 2010.

[12] Y. Chuang, P. M. Smith, L. E. Moser, and I. M. Lombera. Protecting the iTrust Information Retrieval Network against Malicious Attacks, Journal of Computing Science and Engineering, vol. 6, no. 3, Sep. 2012.

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

[14] RIPE DB. ftp://ftp.ripe.net/ripe/dbase.

[15] RADb. ftp://ftp.radb.net/radb/dbase.

[16] CAIDA prefix origins. http://www.caida.org/data/routing/routeviews-prefix2as.xml.

[17] RouteViews Project. http://www.routeviews.org.

[18] RIPE NCC Routing Information Service. http://www.ripe.net/data-tools/stats/ris/routing-information-service.

[19] UCLA IRL Topology. http://irl.cs.ucla.edu/topology.

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

[21] CAIDA Ark. http://www.caida.org/projects/ark.

[22] H. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. Anderson, A. Krishnamurthy, and A. Venkataramani. iPlane: An Information Plane for Distributed Services. In *USENIX NSDI*, Nov. 2006.

[23] 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 JSAC*, vol. 29, no. 9, Oct. 2011.

[24] 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.

[25] NANOG. http://www.nanog.org.

[26] Arin OriginAS data. ftp://ftp.arin.net/pub/originAS.
[27] Tier-1 ASes.
http://en.wikipedia.org/wiki/Tier1network.
[28] AfriNIC. http://www.afrinic.net.
[29] APNIC. http://www.apnic.net.
[30] ARIN. http://www.arin.net.
[31] LACNIC. http://www.lacnic.net.
[32] RIPE NCC. http://www.ripe.net.