

IP Reachability Differences: Myths and Realities

He Yan

Colorado State University
yanhe@cs.colostate.edu

Benjamin Say

Colorado State University
bsay@cs.colostate.edu

Brendan Sheridan

Colorado State University
sheridan@cs.colostate.edu

Dave Oko

Colorado State University
dave.oko@gmail

Christos Papadopoulos

Colorado State University
christos@cs.colostate.edu

Dan Pei

AT&T Research
peidan@research.att.com

Dan Massey

Colorado State University
massey@cs.colostate.edu

Abstract—Ideally, an Internet Service Provider offers reachability to the entire Internet. However, it is often claimed (supported by anecdotal evidence) that reachability differences exist between ISPs, with the result that some valid, globally routable IP addresses can be reached from one ISP but not from others. Outdated Bogon filters, specialized services, and regional differences have all been conjectured as possible sources of reachability differences. However, to the best of knowledge, no systematic study has verified and quantified the extent of such differences, and no analysis exists to show which (if any) of the common conjectures are the dominant causes. In this paper we provide an initial systematic study of top-tier ISP reachability differences. We compare global BGP tables from 25 ISPs to 1) quantify the extent of reachability differences, and 2) determine which (if any) of the common conjectures explain these differences.

I. INTRODUCTION

Internet Service Providers (ISPs) provide their customers with IP reachability. Ideally, a customer purchasing connectivity from ISP A would enjoy reachability to every valid IP address in the Internet. However, there is anecdotal evidence that in practice there are address ranges that ISP A can reach but ISP B cannot (and vice versa). In discussions with both researchers and operators there is widespread consensus that differences do indeed exist. However, to date there has been no systematic study to quantify the extent of these differences. Thus, our first objective in this paper is to *demonstrate that reachability differences do exist and quantify the extent of these differences*. Our study uses a sample of 25 upper-tier ISPs.

A number of conjectures have been put forward to explain reachability differences. One often repeated conjecture is that *Bogon* filtering causes reachability differences for newly allocated prefixes. According to the RIPE report on De-Bogonising New Address Blocks (Bogon filtering) [6], ISPs are increasingly filtering unallocated address space. ISPs less diligent about maintaining their filter tables will wrongly filter out newly allocated and announced prefixes. Thus, different policies in Bogon filtering at ISPs A and B will result in different reachability when one ISP filters out a prefix and another ISP does not. Current data tracking of Bogon visibility can be found at [3].

Another conjecture for reachability differences is that an ISP may advertise services only to its own customers. For example, an ISP may provide caching DNS servers to its

customers. However, for management and security reasons the DNS caches may not be available to non-customers. As result, there may exist prefix ranges which ISP A announces only to its customers which are not available to ISP B . ISP B may follow the same practice, thus resulting in different address ranges visible from ISP A and ISP B .

Yet another conjecture is that some regions of the internet simply have poor global connectivity. ISP A may offer routes to prefixes in say, a developing country, but ISP B does not provide connectivity to the same region. The reasons for these differences may be political, economic, technical, or a combination of the above. The net result is that ISP A and ISP B do not provide the same reachability.

The above list of conjectures is not intended to be complete, but are clearly worthy of additional investigation as they could produce reachability differences. Measurements have shown that Bogon filtering policies do impact reachability [3]. However, the data in [3] does not show whether Bogon prefixes are in fact the accountable for most of the observed reachability differences (in fact, our data shows they are not). Thus, our second objective in this paper is to *analyze reachability differences and show which conjectures (if any) provide a significant contribution to the differences*.

While our results are not likely to solve reachability differences, we believe that they are useful and informative to the networking community, both in research and in operations.

II. COMPARING ISP BGP ROUTING TABLES

We obtained BGP table data from Oregon RouteViews [2] during a 20-day period from April 1st, 2010 to April 20th, 2010. We selected 25 ISPs, which consistently announced their full BGP routing tables (defined as significantly fewer than 310,000 routes - the global routing table size at our study time) to the Oregon RouteViews collector box *routeviews2.routeviews.org*. Table I shows that the average table size for each of these 25 ISP differs, and thus, clearly, the ISPs cannot have identical sets of prefixes (although reachability might still be the same).

To study reachability differences between ISPs, we select ISP0 (a tier-1 ISP where we have access to router configurations and internal BGP tables) as our baseline ISP and compare its BGP table with the BGP tables of 24 other ISPs.

ISP (peering address)	Name	Location	Average Size
ISPO		CA	311310
ISP1 (4.69.184.193)	Level3	SEA	310755
ISP2 (62.72.136.2)	PIPEX	PIPEX	312355
ISP3 (64.71.255.61)	Sprint/Canada	Toronto	313081
ISP4 (66.185.128.1)	AOL	ORD	311491
ISP5 (67.17.82.114)	Global Crossing	Palo Alto	312377
ISP6 (85.114.0.217)	OBITRU	StPetersburg	314658
ISP7 (95.140.80.254)	GBLNETRU	Russia	319244
ISP8 (129.250.0.11)	NTT-A	CA	312008
ISP9 (129.250.0.171)	NTT-A	VA	312008
ISP10 (137.164.16.12)	CENIC	CA	314621
ISP11 (144.228.241.130)	Sprint	Stockton	311438
ISP12 (147.28.7.2)	PSG	SEA	312580
ISP13 (195.22.216.188)	SeaBone	NEW	312088
ISP14 (196.7.106.245)	UUNET	Africa	313653
ISP15 (198.129.33.85)	ESNet	New York	313792
ISP16 (202.232.0.3)	IJJ	Japan	312953
ISP17 (203.181.248.168)	APAN/tppr-tokyo	ORD	317575
ISP18 (203.62.252.186)	Telstra	Sydney	312979
ISP19 (206.24.210.102)	SAVVIS	SF	311431
ISP20 (209.123.12.51)	Net Access	NYC	314335
ISP21 (209.161.175.4)	WCICABLE	OR	313644
ISP22 (213.248.83.252)	Telia	NYC	308592
ISP23 (216.18.31.102)	Group Telecom	BC	312565
ISP24 (216.218.252.164)	Hurricane Electric	SJ	314156

TABLE I
THE AVERAGE BGP TABLE SIZE OF ALL 25 ISPS OVER 20 DAYS.

A. Pair-wise Comparison Methodology

On each day of our study, at approximately the same time, we conducted pair-wise comparisons between ISPO’s routing table and the table for each of the other 24 ISPs. We first identified, and removed, *Exact Matches*. An *Exact Match* is a prefix which appears in the routing table for both ISPO and ISP X . We mark the remaining prefixes as either *Covered* or *Not Covered*. A prefix P that appears in ISP A ’s table is marked as *Covered* if P is covered by a less specific prefix in ISP B ’s table, or if a group of more specific prefixes span P ’s entire range. Otherwise P is *Not Covered*.

Let’s look at a few examples. The prefix “19.0.0/8” in table A is an *Exact Match* if a routing entry for “19.0.0/8” also appears in table B. The prefix “10.1.0/16” in table A is marked as *Covered* if there is a less specific prefix “10.0.0/8” in table B. The prefix “11.2.0/16” in table A would also be marked as *Covered* if the more specific prefixes “11.2.0/17” and “11.2.128.0/17” appear in table B. The prefix “129.82.0/16” in table A would be marked as *Not Covered* if table B does not contain “129.82.0/16”, does not contain a less specific prefix, and *at least some part* of the “129.82.0/16” range is not reachable using more specific prefixes in table B.

B. Pair-wise Comparison Results

Figure 1 shows the average number of prefixes categorized as *Exact Match*. The bars on the graph show the minimum and maximum number of *Exact Match* prefixes. For example, on average 310,500 prefixes appeared in the routing tables for both ISPO and ISP1. Over the 20-day period, there was one day when only 309,500 prefixes appeared in the routing tables for both ISPO and ISP1, and there was another day when nearly 312,000 prefixes appeared in the routing tables for both ISPO

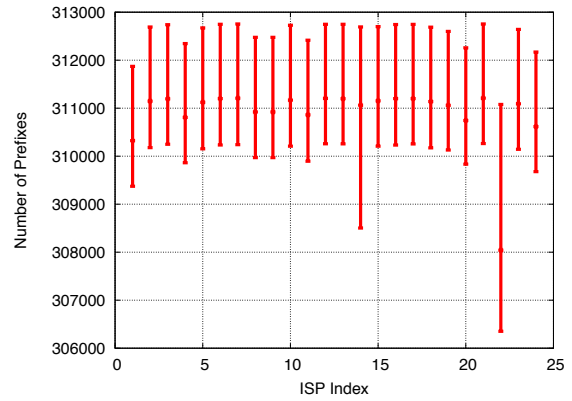


Fig. 1. *Exact Match* prefixes from ISPO’s perspective over 20 days.

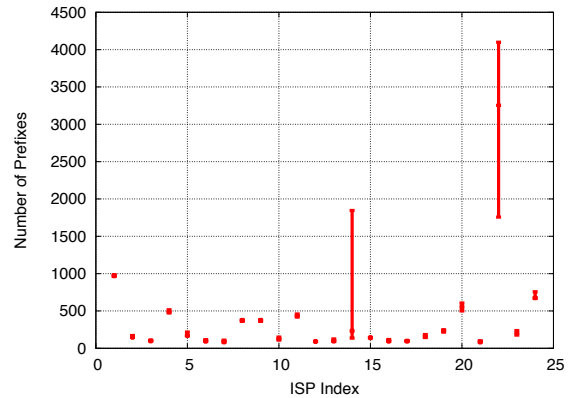


Fig. 2. Non-identical prefixes in ISPO’s BGP tables compared to other ISPs’.

and ISP1. The variations in number of *Exact Match* prefixes largely reflect the dynamic nature of BGP tables.

Given any table from ISPO taken over the 20 day study period, over 99% of prefixes in the ISP’s table were also present at the BGP tables of nearly all other ISPs. One exception is ISP22, where in the worst case only 98.5% of ISPO’s prefixes appeared in ISP22’s table. In other words, nearly every prefix that was present in ISPO’s table was also present in the tables of the other ISPs.

Similarly, the other ISPs had very few prefixes in their tables that were not present in ISPO’s table. The numbers were slightly lower, over 98.5% of prefixes in an ISP table were also present in ISPO’s table. Again there were small exceptions: only 97.25% of the prefixes in ISP7’s table were present in ISPO’s table, and only 98% of the prefixes in ISP17’s table were also present in ISPO’s table.

Figure 2 shows the average number of prefixes which appear in ISPO’s table but do not appear in the corresponding ISP’s table, after removing the *Exact Match* prefixes. The bars indicate the maximum and minimum values observed during the 20-day study. In most cases, the number of non-identical prefixes between ISPO’s table and each other ISP’s table is less than 500.

Figure 3 considers whether prefixes that appeared in ISPO’s table are *Covered* by a less specific (or set of more specifics)

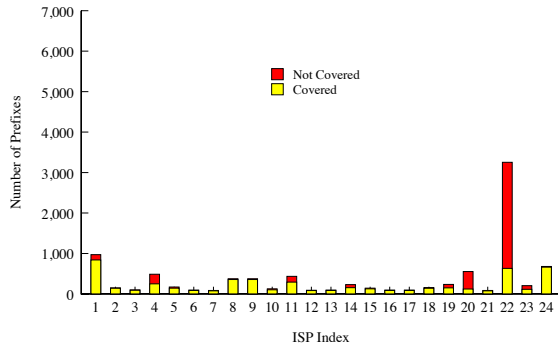


Fig. 3. ISP0’s “Covered” and “Not Covered” prefixes.

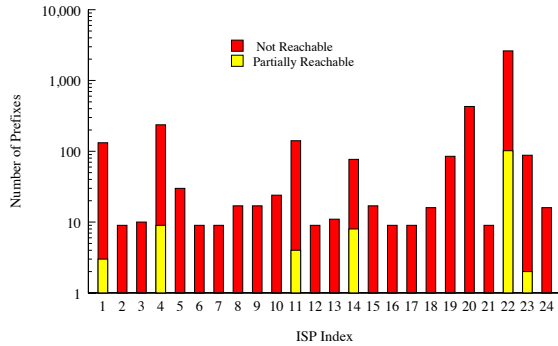


Fig. 4. ISP0’s *Partially Reachable* and *Not Reachable* prefixes.

in the other ISPs’ tables. With the exception of ISP22, which we will explain later, nearly every prefix that appears only in ISP0’s table is *Covered* by some prefix in the other ISPs’ tables.

The overall results show that there are differences between the BGP routing tables, but the vast majority of prefixes appears in both ISP0’s table and the tables of all other ISPs. When there are differences, the prefixes appearing in only one table are typically *Covered* by prefixes in the other table. In terms of BGP routing table reachability, there is in fact relatively little difference between all 25 ISPs in our study.

III. CHARACTERIZING REACHABILITY DIFFERENCES

We now consider the *Not Covered* prefixes only. These prefixes represent differences in reachability as viewed from the perspective of the BGP routing table. To understand the reachability impact we must further divide the *Not Covered* prefixes into two groups. A prefix is *Not Reachable* if none of the IP addresses under a prefix in table *A* can be reached by any prefix in table *B*. A prefix is *Partially Reachable* if some of the IP addresses under the prefix in table *A* can be reached by more specific prefixes in table *B*. For example, *Not Covered* prefix “10.1.0.0/16” in table *A* is marked as *Partially Reachable* if a more specific prefix such as “10.1.0.0/17” appears in table *B*. Figure 4 shows a breakdown of the the *Partially Reachable* and *Not Reachable* prefixes from ISP0’s perspective. Note that the Y axis is in log scale.

We also measured the impact of a *Partially Reachable* or *Not Reachable* prefix in terms of the number of affected IP

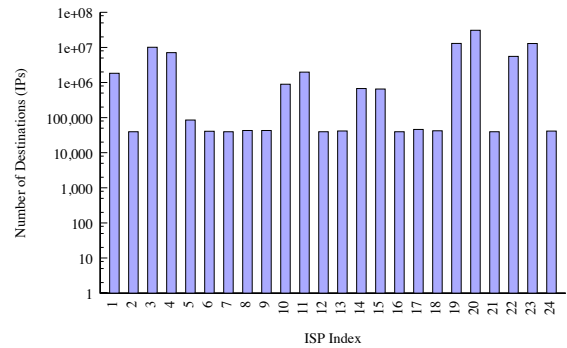


Fig. 5. ISP0’s average number of unreachable IPs in *Not Covered* prefixes.

addresses. For a *Not Reachable* prefix, the number of affected IP addresses is simply a function of the prefix mask.

For a *Partially Reachable* prefix, we first perform the same procedure as above and then subtract the number of IP addresses reachable via some more specific prefixes. Figure 5 shows the number of IP addresses which can be reached by ISP0 but cannot be reached by the other 24 ISPs. The Y-axis uses a log scale. The large number (tens of millions) of destinations that are not reachable from some of the other 24 ISPs can be explained by a few “/8” prefixes that only appear in ISP0’s BGP tables.

A. Temporal Dynamics

BGP routing tables are dynamic and thus a prefix may be *Not Covered* at one instant and then change to either *Covered* or “ExactMatch” after a routing table change. A prefix could also be simply withdrawn from the routing tables and thus cease to be *Not Covered*. Next, for each *Not Covered* prefix we calculate the number of days it appeared during our 20-day study period.

Figure 6 shows the distribution of the number of days prefixes appeared as a *Not Covered* prefix from the perspective of ISP0. The figure shows three large groups at one day, 13 days, and 20 days. The prefixes appearing for one day appear to be a group of transient differences, where ISP0 and at least one other ISP could not reach the prefix for one day. During the other 19 days, either ISP0 no longer announced a route to that prefix, or all 24 other ISPs announced a route to the prefix. Given the dynamic nature of the BGP routing table and its well-known convergence delays, we regard such transient behavior normal and we did not explore it further [9].

On the other hand, the prefixes appearing as *Not Covered* for 20 days indicate that ISP0 could reach these prefixes during every day of our study; moreover, at least one ISP could not reach these prefixes on every day of our study. We investigate this behavior further in the next section. We also show that all but 16 of the prefixes in the 13 day group can be attributed to a single ISP and are advertised over a single geographic connection.

IV. EXPLAINING REACHABILITY DIFFERENCES

So far we focused on the long lasting reachability differences that were identified in the previous sections. Here, we are specifically interested in the prefixes which were




Fig. 6. ISPO's number of *Not Covered* prefixes.

Not Covered for the entire 20 days of the study as seen in Figure 6. These 1750 unique prefixes were present in ISPO's routing table for the entire study period. Furthermore, for each day of the study these prefixes were *Not Reachable* or only *Partially Reachable* by at least one of the 24 other ISPs. Could conjectures such as Bogon lists or regional impacts explain the differences?

A. ISP Groupings

We begin by examining which of the other ISPs had difficulty reaching these 1750 unique prefixes. For each prefix, we calculate which ISPs do not cover the prefix on day 1. It could be the case that all 24 ISPs, or only one of the ISPs do not cover the prefix on day 1. Note that, by definition, each one of these prefixes is *Not Reachable* or *Partially Reachable* by at least one of the 24 ISPs on day 1. This process produced groups of ISPs for each prefix. For example, suppose prefix 10.0.0.0/8 was *Not Covered* by ISP 5 and was either an *Exact Match* or *Covered* by all other ISPs. This prefix produces the group *ISP5*. If ISP7 also did not cover the prefix, the group would be *ISP5, ISP7*. We then repeated the process for each subsequent day of the study. Over 20 days, the 1750 prefixes produced only 28 distinct groups shown in Table II. For all but 71 of the 1750 prefixes, the prefixes produced the same group on each day of the study. The 71 prefixes in the last two rows of Table II had some short-lived differences, where one ISP did not cover the prefix for all days and the other ISP only did not cover the prefix for between 1-7 days.

B. Customer Only Prefixes

Customer Only Prefixes have been suggested as one explanation for ISP reachability differences. As the name suggests, these prefixes are announced to customers but not announced to peers (or providers). None of the ISPs in our study are customers of each other, thus we would expect customer-only prefixes to appear in ISPO's BGP table and to not appear in the BGP tables of any other ISP. Table II (group 26) contains a total of 9 such prefixes. We confirmed that all 9 prefixes were indeed ISPO customer-only prefixes. We also learned that ISPO has substantially more than 9 customer-only prefixes. None of these other customer-only prefixes appeared as *Not Covered* in our study. Instead, they were covered by less specific prefixes that did appear in the other tables. Clearly, the conjecture that

Group #	Group Size	Prefixes Producing this group	ISPs in Group
1	1	1126	ISP22
2	1	288	20
3	1	41	4
4	1	14	5
5	1	8	11
6	1	7	23
7	1	6	14
8	1	4	1
9	1	3	10
10	1	3	19
11	1	3	24
12	1	2	15
13	2	10	20, 22
14	2	2	8, 9
15	2	1	11, 15
16	2	1	22, 23
17	2	1	5, 19
18	2	1	5, 24
19	3	102	1, 4, 11
20	3	37	19, 20, 23
21	3	4	1, 11, 19
22	3	3	1, 4, 19
23	3	1	15, 20, 22
24	4	1	5, 14, 15, 22
25	7	1	1, 4, 11, 15, 19, 20, 23
26	24	9	ALL ISPs
27	2	70	4, 14
28	2	1	15, 22

TABLE II
ISP GROUPS FOR *Not Covered* PREFIXES

customer-only prefixes account for reachability differences is correct, but it explains only 9 of 1750 long-lasting reachability differences in our study.

Customer-only prefixes can completely explain the difference with many ISPs. After removing the 9 customer-only prefixes, ISP2, ISP3, ISP6, ISP7, ISP12, ISP13, ISP16, ISP17, ISP18 and ISP21 do not appear in any group (see Table II). After discounting the customer only prefixes, these ISPs are not responsible for *Not Covered* prefixes that were present for all 20 days of the study. After removing the customer only prefixes, these ISPs covered every prefix that appeared in ISPO's table during the 20 day study.

C. Prefix and Peer Location

Another conjecture to explain reachability differences is that some regions, and perhaps developing regions in particular, are less likely to have global reachability. Thus, ISPs in some regions may have fewer prefixes (or simply different prefixes) when compared to ISPs in other regions. To investigate this conjecture we performed standard whois queries for all 1750 prefixes. By default, whois results are routed to the relevant server for the organization to whom IANA made the regional /8 assignment. 174 of the prefixes yielded no whois results. We grouped the remaining 1,576 based on the country code of the "country:" field of the whois results.

Grouping prefixes by country as described, shows that Bulgaria (BG) was listed for 509 of 1,576 prefixes, making it significantly more common than any other country. However, upon closer inspection, all 509 of those prefixes were consistently unreachable from ISP 22 only. On the other hand, there was no discernible pattern between prefixes with other country codes and or any other ISP. In addition, Great Britain

is the second most common country with 168 occurrences. We conclude that although patterns in reachability based on region do exist, they were uncommon in our study and typically linked to a single ISP. There was also no correlation between developing regions and *Not Covered* prefixes in our data set.

Our data reveals a clear distinction with respect to ISP 22. Table II (Group 1) shows that this ISP was solely responsible for 1126 of the 1750 prefixes. As stated above, 509 of these 1126 prefixes had BG as the country code. We did not find any other patterns to the remaining prefixes.

D. Revisiting the 13-Days Spike

so far, this section has focused entirely on prefixes that were listed as *Not Covered* for all 20 days of the study. Figure 6 shows nearly all *Not Covered* prefixes were either transient (occurring during only 1 day) or long lasting (occurring over the entire 20 day period). The one exception is the notable 13 day spike. There are 1906 unique prefixes that made up the 13-day spike. In the process of examining ISP22 in the section above, we found that 1892 of these prefixes were associated with ISP22. In other words, every ISP, with the exception of ISP22, does cover these prefixes.

We also performed whois queries for these 1892 prefixes to look for any additional commonalities. Notably, 1845 of them yielded country codes for India in the "country" section of the whois response. The prefixes fell into the 13-day bar since ISP22 added all 1892 prefixes to its table on day 14 and maintained those prefixes for the remainder of the study period.

E. Bogon Prefixes and RADb

Bogons prefixes [5] are typically defined as IP addresses which have either not been assigned by IANA or have been marked as reserved or private via RFC1918 and RFC5735. These addresses are regularly blocked by ISPs because traffic from them is typically illegitimate. However, because Bogon lists regularly change over time due to new prefix assignments by IANA, it is not uncommon for an out-of-date Bogon list to block legitimate traffic [6]. New prefixes do experience reachability issues and this can explain some reachability differences between ISPs. While Bogon prefixes can be a factor, it is not clear from the previous work that Bogons are a primary cause of ISP reachability differences.

To better understand the impact of Bogon lists in our study, we first looked to see how many prefixes had been recently removed from Bogon lists. Using the current IPv4 Bogon list and changelog provided by Team Cymru [5], we reconstructed Bogon lists for time between the end of our study and February 13, 2003 (the usable beginning of the changelog). For each period of time, we checked how many prefixes of the group of 1750 with structural reachability differences fell under prefixes in the Bogon list; the results are displayed in Figure 7.

Going back 7 months from the beginning of our study, only 41 prefixes fell under the Bogon list. However, an additional 295 prefixes from the study fell under the 175/8 and 182/8 prefixes assigned to APNIC in August 2009. After that, we saw a steady increase in addresses falling under the Bogon list correlating to periodic assignments made by IANA. There was

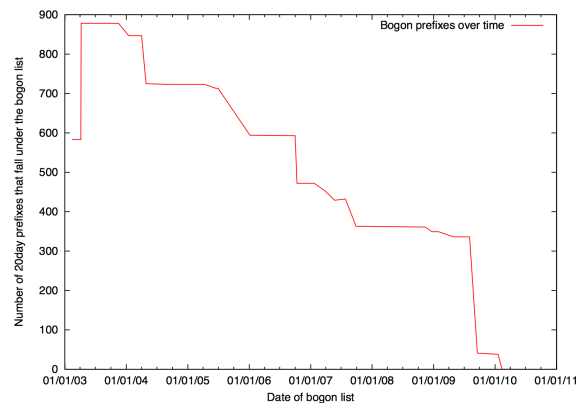


Fig. 7. 20-day Prefixes on the Bogon list over time

also a notable drop in April 2003 related to the new reservation of 17 /8 prefixes. We conclude that while out-of-date Bogon lists likely accounted for some portion of prefixes in our study, they certainly did not account for a majority.

In addition to filtering Bogon prefixes, some ISPs use the RADb database [8] to filter BGP announcements, which is conjectured to lead to reachability differences. To investigate this claim, we ran whois queries against the RADb database (whois.radb.net) for all of our 1750 prefixes with reachability differences. We found that 92 prefixes (roughly 5%) were not present in the database at all, and that the remaining 1,658 yielded 2,472 route entries. Of those 2,472 route entries, 199 (roughly 8%) returned a warning that: "This route object is for a [ISP NAME] customer route which is being exported under this origin AS. This route object was created because no existing route object with the same origin was found, and since some [ISP NAME] peers filter based on these objects this route may be rejected if this object is not created."

We also ran whois queries against the RADb database for all 332,614 routes in the study and found that 38,871 (over 11%) were also not present in the database at all, and the remaining 293,743 yielded a total of 558,301 route entries. Of those 558,301 route entries, 110,481 (nearly 20%) produced the same warning. Given that the occurrences of null entries and warnings in the RADb database were actually more common in routes without reachability differences, we conclude that while reachability differences based on RADb filtering are possible, there is no correlation to support the argument that they account for a significant portion of reachability differences we observed.

F. Prefix Usage

Another conjecture is that the prefixes with reachability differences belong to specialized networks such as those associated with exchange points, test networks, and otherwise inactive networks.

To determine if the prefixes that persisted throughout all 20 days of the study were in use and hosting services, we scanned these networks for some well known ports: 20-30, 53, 80, 110, 443, 993, 995, and 8080 using nmap. To speed up the scanning we allowed nmap to use host discovery, its default behavior, before scanning the aforementioned ports.

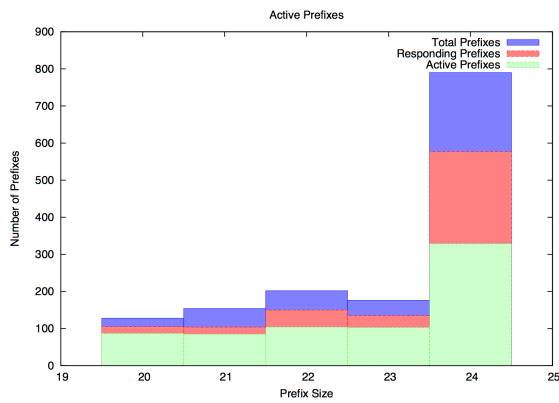


Fig. 8. Breakdown by prefix mask.

Due to both time limitations and security concerns, we only scanned /19 or more specific prefixes. This accounted for 1450 of the 1750 prefixes. Of the 1450 prefixes scanned, 1071 had running services. We then verified that hosts offering services in these prefixes also had DNS entries. This resulted in 708 prefixes (termed as *active prefixes*) with both running services and corresponding DNS entries. The bars in Figure 8 show the total prefixes scanned, the number of prefixes that responded to host discovery, and the active prefixes.

Our analysis suggests that most prefixes support a number of services. In an ad hoc exploration of these prefixes, we found that these active prefixes contained, the following: the Royal Burnie Airlines (202.12.26.120), e-dot, a Hungarian registrar and web hosting company (95.131.48.19), as well as 29 Turkish universities.

Using standard whois queries, we obtained email addresses for 1093 of the 1750 prefixes and contacted each address. The vast majority received no response or automated responses. However, 16 respondents reported no prior reachability issues, 7 reported that the prefixes were in the process of being removed from Bogon filters, 7 reported to have had connectivity or reachability issues but were unsure of the cause, 8 reported the prefixes were undergoing maintenance, being provisioned, or not yet in production, 2 reported being blacklisted by some ISPs due to spam issues, 2 were supposed to be off-line and were not being legitimately used, and finally 1 was only intended to be announced locally.

V. RELATED WORK

This study focused on the reachability differences seen in BGP routing tables. We examined BGP tables from several ISPs and identified which prefixes are reachable via the BGP table in one ISP, but not reachable via any BGP route in other ISPs. There are several popular conjectures about the causes of reachability differences between ISPs. However, to the best of our knowledge, there has not been a systematic comparison of actual BGP routing tables. In this sense, our work is the first to investigate this area and confirm or contradict common conjectures about ISP reachability differences.

It is important to note this is a control plane study, where we focus on quantifying differences in ISP BGP routing tables. In

[4], the authors compared control and data plane and showed that in some occasions control plane reachability does not guarantee data plane reachability (and vice versa). In this work, we focus on the control plane and do not investigate cases where control and data reachability diverge. Interested readers should review [4] for a more complete discussion of control plane/data plane differences.

Other related work on IP reachability has explored how to define, model, and estimate IP reachability from a particular location [7], [10], [11], which can be very useful for troubleshooting, auditing, management and security. These papers observe that complex factors can influence whether a particular IP is in fact reachable from a given location. For example, a BGP route may exist for the IP address but TOS bits in the IP header may make the route unusable for some packets. NAT, firewalls, router ACLs, middle boxes and so forth, all contribute to make the problem of determining IP reachability a challenging task as it is very hard to define, model and measure. Our work is complimentary, because while the studies above do provide some models and techniques to calculate reachability from a given site, they do not address the broader problem of reachability differences among ISPs.

VI. CONCLUSION AND FUTURE WORK

In this paper we compared BGP routing tables from a number of ISPs and investigated BGP level reachability differences between them. Although it is widely assumed that there are reachability differences between ISPs, our results are the first to quantify the differences. We also explored common conjectures for these differences and found that while instances of each conjecture did hold, none of them alone was sufficient to explain even a majority of differences. Finally, we confirmed that reachability differences do affect ISPs who intend to provide a globally reachable service.

REFERENCES

- [1] RIPE routing information service project. <http://www.ripe.net/>.
- [2] RouteViews routing table archive. <http://www.routeviews.org/>.
- [3] Ripe routing information service - bogon tracking. <http://www.ris.ripe.net/debogon/>, 2010.
- [4] R. Bush, O. Maennel, M. Roughan, and S. Uhlig. Internet optimometry: assessing the broken glasses in internet reachability. In *IMC '09: Proceedings of the 9th ACM SIGCOMM conference on Internet measurement conference*, pages 242–253, New York, NY, USA, 2009. ACM.
- [5] T. Cymru. The bogon reference. <http://www.team-cymru.org/Services/Bogons/>, 2010.
- [6] D. Karrenberg. De-bogonising new address blocks. *RIPE Document 351*, 2005.
- [7] A. Khakpour and A. Liu. Quantifying and querying network reachability. In *ICDCS '10: In Proceedings of the 30th International Conference on Distributed Computing Systems*, 2010.
- [8] M. Networks. Radbl the routing assests database. <http://www.radb.net/>.
- [9] D. Pei, B. Zhang, D. Massey, and L. Zhang. An analysis of path-vector routing protocol convergence algorithms. *Computer Networks*, 50(3):398–421, 2006.
- [10] G. G. Xie, J. Zhan, D. A. Maltz, H. Zhang, A. Greenberg, G. Hjalmtysson, and J. Rexford. On static reachability analysis of ip networks. In *in Proc. IEEE INFOCOM*, 2005.
- [11] B. Zhang, T. E. Ng, and G. Wang. Reachability monitoring and verification in enterprise networks. In *SIGCOMM '08 Poster: ACM SIGCOMM Poster Session*, 2008.