



ELSEVIER

Available online at www.sciencedirect.com

SCIENCE @ DIRECT®

Computer Networks 44 (2004) 737–755

COMPUTER
NETWORKS

www.elsevier.com/locate/comnet

Towards capturing representative AS-level Internet topologies

Hyunseok Chang ^{a,*}, Ramesh Govindan ^b, Sugih Jamin ^a, Scott J. Shenker ^c,
Walter Willinger ^d

^a Department of EECS, University of Michigan, 1301 Beal Ave., Ann Arbor, MI 48109-2122, USA

^b Department of CS, University of Southern California, Los Angeles, CA 90089-0781, USA

^c ICSI, 1947 Center St., Suite 600, Berkeley, CA 94704-1198, USA

^d AT&T Labs-Research, 180 Park Ave., Florham Park, NJ 07932-0971, USA

Received 18 June 2003; received in revised form 2 November 2003; accepted 3 November 2003

Responsible Editor: J. Crowcroft

Abstract

Recent studies on AS-level Internet connectivity have attracted considerable attention. These studies have exclusively relied on BGP data from the Oregon route-views [University of Oregon Route Views Project, <http://www.routeviews.org>] to derive some unexpected and intriguing results. The Oregon route-views data sets reflect AS peering relationships, as reported by BGP, seen from a handful of vantage points in the global Internet. The possibility that these data sets may provide only a very sketchy picture of the complete inter-AS connectivity of the Internet has received little scrutiny. By augmenting the Oregon route-views data with BGP summary information from a large number of Internet *Looking Glass* sites and with routing policy information from Internet Routing Registry (IRR) databases, we find that (1) a significant number of existing AS peering relationships remain hidden from most BGP routing tables, (2) the AS peering relationships with tier-1 ASs are in general more easily observed than those with non-tier-1 ASs, and (3) there are at least about 40% more AS peering relationships in the Internet than commonly-used BGP-derived AS maps reveal (but only about 4% more ASs). These findings point out the need for continuously questioning the applicability and completeness of data sets at hand when establishing the generality of any particular Internet-specific observation and for assessing its (in)sensitivity to deficiencies in the measurements.

© 2003 Elsevier B.V. All rights reserved.

Keywords: Internet topology; BGP routing tables

1. Introduction

In the two years prior to the work reported in this paper, ¹ there has been a significant increase

* Corresponding author. Tel.: +1-734-936-0393; fax: +1-734-763-8094.

E-mail addresses: hschang@eecs.umich.edu (H. Chang), ramesh@usc.edu (R. Govindan), jamin@eecs.umich.edu (S. Jamin), shenker@icir.org (S.J. Shenker), walter@research.att.com (W. Willinger).

¹ Most of the work reported in this paper was done in 2001, but the main findings have been checked against the data sets collected in 2003 and continue to hold.

in research activities related to studying and modeling the Internet topology, especially at the level of *autonomous systems* (ASs). For example, these activities include inferring the Internet's AS connectivity graph to describe its properties [2], explaining the origins and causes of some of the observed surprising features [3,4], building topology generators that produce random graphs that resemble the measured AS connectivity graph [5–7], investigating the problem of routing path inflation [8–10], studying the effectiveness of proposed algorithms for detection/prevention of attacks on the network infrastructure [11], and evaluating the performance of multicast protocols [12]. A closer look at the measurements that form the basis for all these studies reveals that the data sets used consist of BGP routing tables collected by the Oregon route server [1]. The Oregon route server connects to several operational routers belonging to commercial ISPs solely for the purpose of collecting their BGP routing tables—we call this data set the *Oregon route-views*. From November 1997 to March 2001, the Oregon route-views have been archived on a daily basis by the National Laboratory for Applied Network Research (NLNR) [13]. Presently, archives of the Oregon route-views are available from routeviews.org [1]. In addition to the full BGP routing table snapshots, routeviews.org also provides daily archives of individual route updates obtained from the Oregon route server.

By making these data sets available to the public, the Oregon route server, the participating ISPs, and the archival sites are providing invaluable service to the research community. ISPs are generally reluctant to disclose information regarding their peering relationships and routing policies. Consequently, the existence of the Oregon route-views data sets has been crucial for enabling and driving the recent research activities on AS-level Internet connectivity. The dearth of available measurement data aside, the use by researchers of these Oregon-based data sets for the purposes of studying the Internet's AS connectivity structure raises the following important issue. The ability to

infer AS peering relationship² from BGP routing tables depends largely on inter-AS business contracts. If a business contract does not permit a given inter-AS route to be used by a third party, BGP does not advertise this information to the global Internet. Additionally, since BGP is a path-vector protocol [14], backup links connecting multi-homed ASs may not show up in BGP routing table snapshots. Consequently, BGP-derived AS connectivity data may yield a very incomplete picture of the actual AS-level Internet connectivity. The authors of [15] raise the possibility that BGP-derived AS-level topology snapshots may not be complete and that extracting path information from BGP route updates may be a better method for obtaining more complete AS topologies. More recently, the router-level connectivity study in [16] suggests that currently available BGP routing tables may not capture many existing AS peering relationships. These papers do not attempt, however, to quantify the extent to which the AS topology information derived from BGP table snapshots may be incomplete.

There has been anecdotal evidence and an intuitive understanding among researchers in the field that BGP-based AS-level topology is not complete. However, as far as we know, there has been no systematic study on *quantifying* the completeness of BGP-derived AS-level topologies. One of the main contributions of this paper is to develop a methodology that enables quantitative investigations into issues related to the (in)completeness of BGP-derived AS maps. Our methodology is as follows. We augment the Oregon route-views with (1) full BGP table dumps from a dozen additional public route servers, (2) a selection of Internet *Looking Glass* sites that provide BGP summary information, and (3) the Internet Routing Registry (IRR). By processing the available BGP dumps, we end up with about 40 BGP

² In this paper, we will use the term *AS peering relationship* to mean that there is “at least one direct router-level connection” between two existing ASs, and that these two ASs agree to exchange traffic by enabling BGP between them. “Provider-consumer” relationship or “peer-to-peer” relationship refer to the *contractual* characteristics of a given AS peering relationship.

views (defined in Section 2), all originating from different ASs. This BGP-derived connectivity data allows us to explore the question of how well the peering relationships maintained by a given AS (“local view”) are observed by other ASs (“non-local view”). We find that a significant number of existing AS peering relationships, especially those among non-tier-1 ASs (defined in Section 2.3), are commonly hidden from most BGP views. We also observe that this phenomenon can be intuitively explained by existing inter-AS peering relationships. In short, these findings reaffirm our earlier comment on the delicate aspects of using BGP data for the purpose of AS-level topology discovery, and suggest that the actual Internet maintains a much richer connectivity structure at the AS level than has been previously reported.

To quantify the difference between the BGP-derived AS connectivity and the actual inter-AS peering relationships, we consult IRR databases that maintain individual ISP’s routing policy information. The IRR’s goal in maintaining these databases is to coordinate and facilitate the setting of global routing policies. Considering potential inaccuracy associated with such manually maintained datasets, we carefully sanity-check the IRR information before using it in our study. We find that AS graphs reconstructed from the Oregon route-views data sets, the Looking Glass sites, as well as IRR information have typically about 40% more edges (and about 4% more nodes) than their counterparts that rely solely on the Oregon route-views data.

The implications of our findings are twofold. First, they clearly demonstrate the need for heightened awareness of, and criticality towards, relying on any single data repository. Even when the data is by itself of the highest overall quality, its applicability and sufficiency should be evaluated in terms of the particular needs of any given study. For example, many of the reported results about routing path inflation, the effectiveness of algorithmic solutions to network security problems, or performance comparisons of different proposed protocols could be strengthened by examining their (in)sensitivity to incomplete connectivity information.

Second, as far as published AS connectivity studies are concerned, our findings have practical as well as theoretical implications. For example, the finding reported in [2] claiming that measured AS graphs exhibit power-law vertex degree distributions, can be interpreted qualitatively to mean simply that these vertex degrees are highly variable, i.e., they typically vary by over three or so orders of magnitude. This qualitative interpretation is *not* disputed by our findings. However, our findings state that while the vertex degree distributions resulting from more complete snapshots of the AS graph do not conform to the strict power-law characteristics, they are clearly consistent with the more flexible class of heavy-tailed distributions such as the Weibull distribution or the family of distributions where the tail is characterized by a power-law and where the rest of the distribution can be essentially arbitrary.

Clearly, this latter distinction has direct implications for the generation of Internet-like graphs or for the more challenging question of explaining the origins and causes of the highly variable vertex degrees in the Internet context. To illustrate, the work by Barabási and Albert [3,4] takes the quantitative power-law observations at face value and provides a suite of results, including constructions that attempt to explain the causes that lead to power-law vertex degree distributions. The applicability of these results and constructions to the Internet has been claimed in [4], based on the power-laws reported in [2]. Even though the reported constructions can be modified to achieve a better fit to the data and accommodate the observed deviations from the strict powerlaw characteristics (e.g., see [17]), these modifications typically result in more highly-parameterized models—a telling sign that when viewed as a concrete null hypothesis, the proposed model will likely be rejected when validated against relevant measurements (see for example [18]). In turn, our findings motivate the formulation of alternative model candidates that will hopefully be more successful in providing an in-depth physical understanding of the properties of the actual Internet topology at the AS level and of their origins.

The rest of the paper is structured as follows. In Section 2, we introduce the notion of a

representative BGP view and explore in detail how well peering relationships maintained by an individual AS are observed by other ASs. To quantify the degree of incompleteness of BGP-derived AS maps, we include in Section 3 information from the IRR and use the IRR dataset to obtain a more complete picture of the existing AS connectivity (and of which BGP only sees a certain fraction). We conclude in Section 4 by commenting on some of the lessons learned and by highlighting the implications of our findings.

2. On the completeness of BGP-derived AS-level topology

If the actual AS-level Internet topology were known, the completeness of a topology constructed from the Oregon route-views could be checked by comparing it with the actual topology. Since the actual Internet AS-level topology is not known, we adopt the following approach to check the completeness of the topology inferred from the Oregon route-views. The BGP routing table obtained from an AS contains information about that AS's connectivity to other ASs. It also contains information on the connectivity between other ASs. Assume that the BGP routing table collected at an AS X contains the most complete vertex degree information obtainable of AS X .³ The BGP routing table obtained from AS Y will see some, but most likely not all, of the connectivity between AS X and other ASs. Similarly, the BGP routing table obtained from AS Z will see some but not all connectivity between AS X and the other ASs. Taking the union of observations from ASs Y and Z , we will likely get a more complete count of AS X 's vertex degree than from either one of them alone, though by no means the

complete count. Considering that the Oregon route-views are the collection of BGP routing tables obtained from several ASs, the question we ask in this section is, "How many (or possibly which) BGP routing tables from different distinct ASs do we have to aggregate before we see the same vertex degree of AS X as reported by AS X 's BGP routing table?" To answer this question, we first collect BGP routing tables from several distinct ASs.

2.1. Available BGP routing tables

Besides the Oregon route server, the Swiss Network Operators Group (SwiNOG) also provides access to a non-commercial route server that collects and makes publicly available BGP routing table dumps [19]. Additionally, as of April, 2001, 10 commercial Internet Service Providers (ISPs), residing in different ASs, also allow public access to their route servers providing full BGP table dumps. As the very first step of our study, we collected BGP routing tables from all these route servers. Furthermore, we have also obtained address prefixes and AS path information from UUNET. Due to the different nature of the BGP information available at Oregon and SwiNOG from that available at the commercial route servers, we will denote the Oregon and SwiNOG route servers the "collector" route servers, and call the others "operational" route servers. Table 1 lists the characteristics of the route servers. In the table, the Oregon route server is labeled "NC1," the SwiNOG route server "NC2", and the operational route servers "C1" to "C10". The column "# next hops" lists the number of distinct next hop routers found in each BGP routing table, the column "# neighbor ASs" lists the number of distinct ASs those routers reside in.

The commercial route servers connect to other, topologically distributed, internal routers (iBGP routers) residing in the same AS, each of which peers with several external routers (eBGP routers) located in different ASs. Depending on route server configurations, the number of "next hops" reported for commercial route servers is either the number of iBGP routers connected to a given route server (C1–C3 and C5–C9), or the aggregate

³ If BGP-running routers residing within a single AS are configured with slightly different policy routing, which could be the case for ASs with continent-wide geographic scope, the BGP routing table exported by only one of them may not have the complete vertex degree of the given AS. The assumption is not that we have the complete vertex degree but that we have the most complete vertex degree obtainable.

Table 1
BGP dump from public route servers

Name	Operator	AS#	# next hops	# neighbor ASs	Data size (MB)
NC1	Oregon-IX	–	43	35	291.5
NC2	SwiNOG	–	42	16	131.3
C1	AT&T	7018	24	1	145.8
C2	Exodus	3967	199	279	127.0
C3	GT Telecom	6539	7	1	57.1
C4	Global Crossing	3549	3,089	447	49.5
C5	Exodus Europe	8709	19	187	47.4
C6	CERFnet	1740	3	1	26.0
C7	Colt	8220	42	331	20.3
C8	Global Online	4197	12	82	18.3
C9	Tiscali	3257	1	1	8.0
C10	GTE	1	1175	495	7.8

number of eBGP routers seen through iBGP routers (C4 and C10). The number of “neighbor ASs” has to be interpreted in a similar fashion. Finally, when observing the variability in BGP table sizes (“Data size”), we note a lack of correlation between BGP routing table sizes and the number of distinct next hop routers. These observations led us to further scrutinize the data available from each source and to use in our analysis only sources from *BGP viewers* satisfying the criteria below.

2.2. Extracting BGP views

We define the union of all the address space reachable in all the available BGP routing tables as the *known address space*. Next, we define a *BGP viewer* to be either an “operational” route server or a *peer* of a “collector” route server. Ideally, the BGP routing table of a *BGP viewer* must cover the whole known address space. We expect that a given BGP viewer would capture the *complete AS-level connectivity* of its own AS. Given a BGP viewer, we define its *BGP view* as an instance of the AS-level topology constructed from its routing table.

While the routing table of a peer of a “collector” route server contains the whole address space reachable through that peer, the address space reachable through an “operational” route server must be constructed from the routing tables of all its peers. An “operational” route server may see advertisements for a given address space from several of its peers, for instance:

```
* i12.1.245.0/24 193.251.245.72 7018 27532
* i              193.251.245.64 1 27532
*>i             193.251.128.22 7018 27532
* i              193.251.245.47 7018 27532
* i              193.251.129.8 1 27532
```

In this example, the “operational” route server can reach the address space 12.1.245.0 through five of its peers, whose addresses are listed in the second column. The remaining columns list the ASs (the AS path) a packet destined for that address space must travel through, for each of the alternatives. The best AS path for each address prefix, according to the local routing policy set by the administrator of the AS, is marked with a ‘>’ in conventional BGP routing tables. When a BGP router re-advertises a particular route, it advertises only the best path, after prepending its own AS number to the AS path. Therefore, to construct a BGP routing table of an “operational” route server, we use only the best entry for each individual address prefix.

Recall that our goal in this section is to answer the question, “How many BGP routing tables from distinct ASs must we aggregate to capture all the vertex degree reported by the BGP routing table of a given AS?” Our construction of “operational” route servers’ BGP routing table reflects our intention to construct the BGP routing table of an individual AS that can then be used to answer this question. That is, our goal here is not to infer the AS-level topology from the individual

“operational” route servers. If that had been our goal, we would have constructed a BGP routing table consisting of *all* AS paths from all peers of the route servers instead of just the best path. Doing so, however, will only bring us back to our original question of how complete such an AS-map would be. Nevertheless, for completeness sake, we also look at the AS graph constructed from *all* available AS paths later in Section 3 (Table 3).

To summarize, for “collector” route servers, each of their peers is a potential *BGP viewer*; whereas for “operational” route server, we have only a single potential *BGP viewer* whose routing table must be constructed from the routing tables of all its peers. Thus in this study we have 10 potential *BGP viewers* from the 10 “operational” route servers and 85 candidates from the two non-commercial “collector” route servers. In addition, the router from which we obtained the UUNET routing information is also considered a potential *BGP viewer*.

We mentioned earlier that in order to qualify as a BGP viewer, ideally a candidate’s routing table must cover all of the known address space. Practically, since each AS has different prefix filtering policies, the complete known address space may not be visible to all ASs. So rather than requiring the complete coverage of all known address space from BGP viewers, we instead *disqualify* any BGP viewer candidate with relatively limited address space coverage.

To compare the coverage of the address space among our BGP viewer candidates, we look at four different measures in each of the candidates’ routing tables: (1) the number of routes, (2) the number of non-aggregatable routes,⁴ (3) following [20], the number of routes whose prefix length are less than or equal to 24, and (4) the number of origin ASs. Using these four measures, we sort the 96 candidates (i.e., 10 from the “operational” route servers, 85 from the two “collector” route




Fig. 1. Number of routes (origin ASs) vs. rank.

server, plus UUNET routing information) in decreasing order and plot the top 65 in Fig. 1. From the figure, it can be seen that all four measures visibly decrease after the 51st rank or so. It turns out that the four different measures pick out the same set of 51 BGP routing tables. Thus of the 96 candidate BGP viewers, only 51 of them satisfy our definition. As for the remaining 45 candidates with relatively incomplete address space coverage, most of them are peering with our two “collector” route servers. Deliberately or for technical reasons, they don’t provide full BGP feeds to the “collector” route servers. The 51 BGP viewers with more or less complete coverage reside in 41 distinct ASs (for a complete list of the 41 ASs, see [21]). Using the BGP views from these 41 ASs gives us 41 perspectives of the Internet.⁵ All of the 41 BGP views were collected on the same date (25 May 2001) at approximately the same time of day.

2.3. Local vs. non-local BGP view

Given our dataset, we ask, “How well are the peering relationships of a given AS observed by *other* ASs”? For example, can AT&T’s BGP routing tables discover UUNET’s AS neighbors reasonably well? How well will a small ISP’s BGP

⁴ An aggregatable route is a redundant route which could be removed from a given BGP table by route aggregation; e.g., if a BGP table contains two prefixes “12.0.0.0/8” and “12.1.140.0/24” with the same AS path “3786 1 7018”, we say that the route containing “12.1.140.0/24” is redundant in the BGP table.

⁵ When an AS has multiple BGP views, we use one of them in our study. Considering all available multiple BGP views of a given AS does not affect the observations made in Section 2.3.

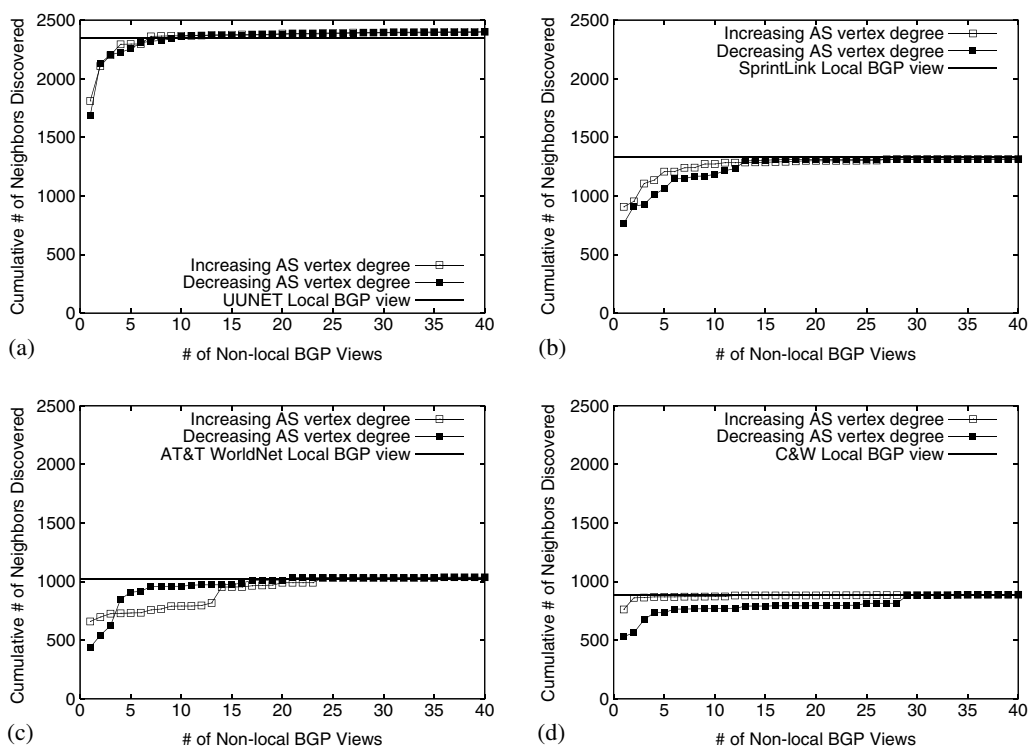


Fig. 2. AS neighbor discovery by non-local BGP views: tier-1 ASs. (a) UUNET (AS701); (b) SprintLink (AS1239); (c) AT&T WorldNet (AS7018) and (d) C&W (AS3561).

view predict AT&T’s AS neighbors? This question has very practical relevance to the goal of our paper since constructing global AS-level topology today has been predicated on collecting a small number of BGP views from the Internet.

To answer the above question, we consider two kinds of BGP views: “local” and “non-local.” From a given AS X ’s perspective, a BGP view originating from an AS X ’s own router is considered *local* and those originating from any other ASs’ routers are *non-local*. Therefore, each of our selected 41 ASs has one local view and 40 non-local views. We assume that any kind of peering relationship maintained by AS X will be best observed in its local BGP view. Based on this, we compare—for each of the 41 ASs—AS X ’s vertex degree predicted from its 40 non-local views against that inferred from its local view. By doing so, we will be able to quantify the completeness of *non-local* views. In this study, we classify the 41 ASs into five hierarchy groups (i.e., tier-1 to tier-5)

as defined in [22], and consider ASs in each hierarchy group separately.

In Figs. 2–4, we look at the number of peering neighbors for a given AS, *cumulatively* discovered by an increasing number of its non-local views. That is, as we incorporate more non-local views (with respect to a given AS), we look at how many more neighbors connecting to that AS are found. The 40 non-local views are merged in two different orders: A non-local view from the highest degree AS is added first, then the non-local view from the second largest AS is added, and so forth (noted as “decreasing AS vertex degree” in the figures). The opposite order (the smallest AS first) is noted as “increasing AS vertex degree”.⁶ The horizontal dotted line in each figure represents the vertex

⁶ The ordering of AS degrees is determined from the union of all 41 BGP views.




Fig. 3. AS neighbor discovery by non-local BGP views: tier-2 ASs. (a) ESnet (AS293); (b) KPNQwest (AS286); (c) Telstra (AS1221) and (d) C&W-Europe (AS12541).

degree predicted from the given AS's local BGP view. In the following, we present our findings for tier-1 ASs and non-tier-1 ASs separately since doing so provides us with more insights into non-local BGP views.

Tier-1 ASs. Fig. 2 shows that a sufficient number of non-local BGP views can discover most of the neighbors connecting to tier-1 ASs.⁷ Interestingly, each non-local view contributes non-uniformly to

the total view. In particular, non-local views from ASs with smaller degrees tend to contribute a larger portion of the total view than ASs with larger degrees, i.e., the curve with the white dots lies above the one with the black dots. This phenomenon can be intuitively explained by the *non-transitive* peer-to-peer relationship and the *transitive* provider-consumer relationship an AS has with its neighbors [23]. The information regarding the pairwise peer-to-peer relationships maintained by a given AS does not circulate among its peers, but does propagate to its downstream customers. Thus, the BGP views from ASs with smaller degrees tend to better observe tier-1 ASs' peering relationships.

Non-tier-1 ASs. For non-tier-1 ASs (Figs. 3 and 4), the combined 40 non-local BGP views clearly fail to observe many existing peering relationships, though there are some exceptions (e.g., Fig. 3(b) and (d)). In these exceptional cases, a single AS (the sixth smallest-degree AS and the smallest-

⁷ In the case of UUNET and AT&T, the 40 merged non-local BGP views find more neighbors than the AS's local view does. We conjecture that for such ASs as UUNET and AT&T, whose geographic presence is continent-wide, a single local BGP view from one location may not be able to capture all their existing neighbors, which are also spread worldwide. Another possibility is that the instability of BGP connections causes some neighbors to be not captured in a given local BGP view snapshot.




Fig. 4. AS neighbor discovery by non-local BGP views: tier-3 ASs. (a) SolNet (AS9044); (b) STARTAP (AS10764); (c) INIT7 (AS13030) and (d) IBS (AS8271).

degree AS, respectively) observes most of the given AS's neighbors. In the case of Fig. 3(b), which is KPNQwest's AS, the sixth smallest-degree AS turns out to be the customer of two other ASs which are in sibling-to-sibling relationship with KPNQwest.⁸ This particular peering relationship allows the sixth AS to receive all the transit routes from KPNQwest, and thus to observe the peering neighbors of KPNQwest well. In the case of C&W-Europe in Fig. 3(d), it turns out that the smallest non-local AS is a regional ISP in Switzerland which is connected to C&W-Europe as a customer. These cases re-confirm how well customer views can discover their

⁸ The peering relationship inference was performed by the heuristics of [23]. Two ASs in sibling-to-sibling relationship are allowed to not only exchange their downstream customer routes, but also export to each other their provider or peer routes.

providers' peering relationships. For a majority of non-tier-1 ASs whose customer views are not available, their peering status is not sufficiently approximated by the merged 40 non-local BGP views. For example, for 14 out of 29 non-tier-1 ASs, more than half of their existing neighbors are hidden from the combined 40 non-local views (see Fig. 6).

Still, it came as a surprise to us that dozens of BGP views of different ASs are hardly sufficient to capture the majority of non-tier-1 AS connectivity. Given that a non-negligible number of an AS's neighbors can be concealed from other ASs, we decided to look more carefully at those missing neighbors.

First, are those missing neighbors caused by hidden *nodes* or hidden *edges*? A *node* is "hidden" if its AS number exists in the local AS's routing table but not in any of the other ASs' routing tables. When an AS is hidden, its address space may still be reachable by other ASs as part of a larger

aggregated address space [24]. In some other cases, such hidden AS numbers may simply be private AS numbers which are used locally by a large ISP to identify subdomains within a given AS [25]. On the other hand, a hidden *edge* means that the neighbor peering relationship between the two end points is not listed in any AS path. Fig. 5 lists, in decreasing order, for each AS *X*, the number of its neighbors not found in any of the other 40 ASs' routing tables. The solid component of each bar is the number of hidden ASs, i.e., ASs whose AS numbers are not present at all in non-local BGP views. The rest is due to undetected peering relationships between AS *X* and its neighbors. One can see that the number of missing neighbors as a component of hidden ASs is negligible (the *y*-axis is in log-scale). The majority of missing neighbors are caused by hidden *edges*, not hidden ASs.

Two recent independently published works [23,22] provide heuristics to infer inter-AS relationships from publicly available BGP data. The heuristics described in [22] classify an AS either as a provider, customer, peer of another, or “unknown.” Additionally, ASs are also classified into five AS hierarchy levels (i.e., from tier-1 to tier-5). We applied the heuristics from [22] to the AS relationships “hidden” from the non-local views and plot the results in Fig. 6. Each bar in the figure corresponds to one of our 41 ASs. The *height* of each bar is the ratio of an AS's hidden neighbors over all its neighbors that are inferred from its local view. For example, some ASs have as much as 80% of their neighbors hidden from their non-local views. Different shaded regions on a bar de-




Fig. 5. Number of missing neighbors.




Fig. 6. AS relationships of missing neighbors.

note the fractions of hidden neighbors who are classified as provider, customer, peer, or “unknown” by the heuristics in [22]. The 41 ASs are further grouped by their positions in the five-level AS hierarchy, proposed in [22]. This grouping is indicated on the *x*-axis (none of the 41 ASs was classified as tier-4 ASs).

In agreement with our earlier analysis illustrated in Fig. 2, Fig. 6 shows that neighbors of the largest ASs located in the tier-1 hierarchy level are reasonably well observed. For some other tier-1 ASs, however, many of the peering relationships with their customers are hidden from non-local views. In case of those tier-1 ASs, we find that almost all their hidden customers are connected to at least one other provider AS (i.e., multi-homed). Similarly, some multi-homed connections of tier-2 and tier-3 ASs are not captured in their non-local views. In BGP, given several available AS paths to each destination, an AS selects one (i.e., the *best*) AS path and only advertises this best path. This feature makes the less preferred upstream connections not visible to the global Internet. Finally, ASs located in lower levels of the Internet hierarchy maintain many peer-to-peer type relationships, which are privately shared by the peers and their customers only, and therefore are not globally available.⁹

⁹ Infrastructures such as public exchange points or network access points allow non-tier-1 ASs to peer with each other at relatively low cost [26].

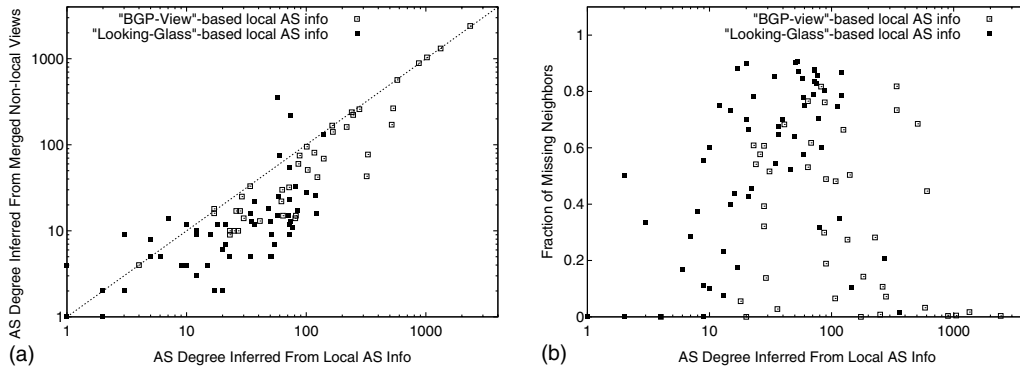


Fig. 7. The completeness of non-local BGP views.

2.4. Other BGP-derived connectivity information

To help troubleshoot Internet-wide routing problems, several ISPs enable public, but limited, access to several of their selected border routers or route servers through the *Looking Glass* tool. By querying an AS's Looking Glass, we can obtain its BGP summary information, i.e., a list of the AS's neighbors and their BGP session status. From this BGP information, we can deduce the set of AS neighbors connected to the local AS.¹⁰

Once we obtain the number of each AS's neighbors from its corresponding BGP summary information, we compare it with the one constructed from non-local BGP views, as in Section 2.3. However, unlike the BGP routing tables of the BGP viewers, BGP summary information from individual Looking Glass sites may not list all their AS's neighbors. Routers from which the summary information is collected may only connect to a subset of all existing eBGP routers. Therefore, the number of neighbors revealed by an individual piece of BGP summary information can only be interpreted as a *partial* view of the AS's neighbors.

In Fig. 7(a), we visualize how well the degree of a given AS inferred by its combined *non-local* BGP views (y-axis) is correlated with its actual degree (x-axis). The actual degree of a given AS is based on either its local BGP view or its Looking Glass data, as labeled. The black dots labeled "BGP-view-based local AS info" summarize the neighbor discovery results of Section 2.3. A dot below the diagonal line means that the degree of the corresponding AS is not well predicted by its non-local BGP views. A dot above the diagonal line means that the local source has a less complete view of the AS degree than the non-local source. Note that all the dots above the diagonal line in Fig. 7(a) are associated with the Looking Glass data. As mentioned earlier, Looking Glass data may not contain the complete neighbor list of a given AS. Aside from these exceptional cases, it is clear that the vertex degrees predicted by non-local BGP views are incomplete.

In Fig. 7(b), we quantify the incompleteness of non-local views illustrated in Fig. 7(a). We first calculate the ratio of an AS's hidden neighbors over all its actual neighbors and then plot the calculated ratio as a function of an AS's actual degree. Similar to Fig. 7(a), the labels within the figure indicate how we infer the actual degree of a given AS. Except for the ASes with very small and very large degrees, most of the ASes in this figure have a significant portion (from 10% to as large as 90%) of their neighbors hidden from their non-local views. The missing ratio is most dramatic for those ASes with degrees between 40 and

¹⁰ Querying a Looking Glass is done through a web-based interface; we pre-selected 60 or so sites [27] and have our crawling script periodically collect their BGP summary information. We started our script at the same time we collected BGP dumps from the different route servers mentioned in Section 2.1.

80 or so, in which case more than half of all their neighbors are concealed from their non-local views.

Contrary to what has been commonly assumed, our study shows that a non-negligible number of existing AS peering relationships can be hidden in most BGP routing tables and that the ability to infer AS peering relationships from BGP routing tables depends to a large extent on the type of inter-AS relationships. This in turn suggests that the Internet might maintain a much richer connectivity than is observed by a handful of BGP routers.

3. Augmenting connectivity using the Internet routing registry

The findings from our BGP-based analysis of the AS-level Internet topology give rise to a more fundamental question: “To what extent does the AS topology derived from the Oregon route-views deviate from the *complete* Internet AS-level topology?”

Our observations imply that to obtain a more accurate picture of the Internet’s AS-level topology, BGP views should be collected from end-customer ASs located in the *lowest* levels of the AS hierarchy. However, we do not know how many such BGP views would be sufficient to discover most of the existing upstream connections. Facing this obstacle, we turn to the Internet Routing Registry (IRR) [28] to further glean local AS connectivity information. The IRR maintains individual ISPs’ routing information in several public repositories in an attempt to coordinate global routing policy. The IRR’s routing policy database stores routing information by using the Routing Policy Specification Language (RPSL) [29]. Individual pieces of routing information expressed in RPSL are called *objects*. The following two hypothetical database objects illustrate how such routing information is expressed using RPSL.

```
route:      1.2.3.0/24
desc:      Foo.com
origin:    AS1
changed:   admin@foo.com 20010313
```

```
source:     RADB

aut-num:    AS1
as-name:    FOO-ASN
desc:      Foo Primary AS
import:     from AS2 action pref=100;
            from AS3 action pref=200;
            accept AS4
export:     to AS2 announce AS4
            to AS3 announce ANY
changed:    admin@foo.com 20010313
source:     RADB
```

The first object states that “1.2.3.0/24” belongs to AS1 as of 13 March 2001. These types of objects, which describe individual address prefixes, are called *route* objects. The latter object, which expresses AS1’s import and export routing policies, indicates that AS1 has two peering neighbors AS2 and AS3 with which it exchanges route reachability information of AS4. These types of objects are called *aut-num* objects. From the import and export policy specification of the AS1’s *aut-num* object, we can infer the neighboring ASs of AS1.

3.1. On the freshness of the Internet routing registry

We next question the reliability of IRR routing policy information which is manually registered and maintained. The motivation of the IRR is to minimize the negative impact of the growing number of ASs and the accompanying complexity of inter-AS connectivity on the Internet routing infrastructure [30]. However, being predicated upon voluntary publication of routing policy, the IRR database may not be complete and some part of it can simply remain out-of-date.

According to [31], an increasing number of ISPs rely on the IRR to filter route announcements at border routers. In particular, the RIPE portion of the IRR is actively used by most ISPs in Europe. Many European exchange points [32–34] specify as a membership requirement that members register their routes and peering policy in the RIPE database, whereas the network access points and exchange points in the US do not require such compliance. Given this requirement, we consider the RIPE database a potentially more reliable source than the rest of the IRR databases. For

example, when comparing the RADB and RIPE databases of 25 May 2001, we found that out of the 2673 ASs registered with RADB, only 2039 (76.3%) published their routing policy; in contrast, 4203 (93.6%) out of the 4492 ASs that had registered with RIPE published their routing policy.

To further compare the freshness of IRR databases, we checked the individual objects of the RADB and the RIPE databases as follows. For each object in the databases, we checked its contents with the daily snapshots of the Oregon route-views data [13,35] which were collected since the object's last update time. Route objects (i.e., registry entries recording specific routes) were checked for their route origin information; aut-num objects (i.e., registry entries specifying the routing policies of ASs) were checked against the ASs' neighbors lists.

When checked with the Oregon route-views data, a given object can be either: (a) consistent, (b) inconsistent, or (c) not verifiable with the data. For a given route object that reports on a prefix P , the object can correctly state that the prefix P originates from AS X as observed in the BGP routing tables (a), or it incorrectly indicates that the prefix P belongs to AS Z (b), or the BGP routing tables simply do not contain the prefix P (c). Similarly, for a given aut-num object that describes AS X , the object can report all its peering relationships shown in BGP routing tables (a), or it fails to report some of them (b), or the BGP routing tables do not show the AS X at all (c).

We downloaded public IRR database files mirrored at [36] on May 25th 2001 and checked each of their objects with the Oregon route-views as described above. We consider an object outdated if its information used to be consistent with an older BGP routing table, but has become inconsistent or not verifiable with the more recent (25th May) BGP table.

Fig. 8 compares RADB and RIPE in terms of freshness. They show the frequency (y -axis) of outdated objects among those that have been last updated within a certain number of days, where the number of days are given on the x -axis. The age of an object thus indicates how many days have passed between the time the object was last updated and 25th May. The figure clearly dem-

onstrates that the RIPE database is maintained more carefully and in a more up-to-date manner than the RADB database.

3.2. Verifying routing registry dataset

Given the observed diversity in completeness and freshness among existing IRR databases, we decided to perform a careful sanity check on available IRR databases before using them in our study. Since our source of obtaining AS connectivity information from IRR databases are aut-num objects, not route objects, we focus on aut-num objects from now on. In the following, we describe two different methods of checking the validity of individual aut-num objects in IRR databases.

The first method is based on identifying three types of invalid aut-num objects:

Void objects. We consider an object *void* if either: (1) the AS described by the object has never appeared itself in the Oregon route-views since the object's last update time, or (2) the AS described by the object was once present in the Oregon route-views (dating from November 1997) but disappeared from the tables afterwards.¹¹

Obsolete objects. To find *obsolete* objects, we first construct an *AS reference graph* from available aut-num objects. The AS reference graph is a directed graph where each node corresponds to a registered AS and an edge corresponds to published peering relationship between an AS and a neighbor. An edge is directed from node A to node B if the aut-num object of AS A specifies AS B as a peering neighbor. If all the objects were up-to-date, all edges in this graph must be bidirectional, since any kind of peering relationship is by definition based on bilateral agreement. A unidirectional edge indicates at least one of the two incident ASs has outdated information in the database. If the update time of the two incident ASs of a unidirectional edge is more than 1 month apart, we consider the object updated earlier obsolete.

¹¹ To prevent artifacts caused by the finite time frame of our data, an AS is considered void only if its disappearance dated more than three months from the date of the last entry of our data set (25 May 2001).




Fig. 8. Frequency of outdated objects: (a) outdated route objects and (b) outdated aut-num objects.

Incomplete objects. We consider an object *incomplete* if the AS described has an existing neighbor which is not registered with IRR. To detect unregistered neighbors of a given AS, we obtain a list of the AS's peering neighbors from available BGP routing tables and Looking Glass data. If at least one of these peering neighbors is found unregistered in the AS's aut-num object, we consider the object incomplete.

Only objects that are neither void, nor obsolete, nor incomplete are considered *valid* objects.

This sanity check, however, can potentially invalidate aut-num objects of larger ASs more easily than those of smaller ASs as it allows any single obsolete or unregistered peering relationship of an AS to nullify the AS's entire aut-num object. Thus, in order to avoid a possible bias when invalidating aut-num objects, we also consider the following less stringent method, which is simply based on the age of the objects. That is, our second method disregards all the aut-num objects whose ages are more than a certain age threshold. We call them *retired* objects. In our experiment, we set the age threshold for retired objects to three months.¹²

Table 2 summarizes the validity analysis of the IRR databases by our two methods described above. With the first method (labeled "Sanity check #1"), we filter out void, obsolete and

incomplete objects in succession, and report in each step the number of remaining objects and the total number of peering relationships found from these objects. Similarly, with the second method (labeled "Sanity check #2"), we report the same two numbers after removing all retired objects.

According to the table, the first method yields only about 25% of existing aut-num objects in the IRR databases as valid information. The age-based second test passes even a smaller percentage (i.e., about 20%) of the IRR datasets as valid. However, in terms of the total number of peering relationships, the second test yields a much larger number of peering relationships than the first test. As alluded earlier, objects registering a large number of peering relationships can easily be caught as outdated by the first rigorous test, but tend to pass the second test.

As Table 3 shows, persuing the IRR database allows us to identify an extra 5000–14,000 or so edges over the most complete AS graph constructed from all available BGP data (compare the last two rows against the third-to-last row of Table 3). The table shows the number of nodes (ASs) and edges contained in the AS graph constructed from the various sources, cumulatively. The first row, labeled "Oregon route-views" lists the number of nodes and edges found in the AS graph constructed from the Oregon route-views. The second row ("+RSs") lists the number of nodes and edges found in the Oregon route views plus the full BGP dumps from 11 public route servers listed in Table 1. Recall that in the BGP view analysis in Section 2, we used only the best paths from each full BGP

¹² According to Fig. 8, the IRR data whose age is less than 100 days or so exhibits about 95% accuracy for both RADB and RIPE.

Table 2
Validity analysis of IRR database

Sanity check #1			Sanity check #2		
Type of objects	# of objects	# of edges	Type of objects	# of objects	# of edges
All objects	7521 (100%)	43,498	All objects	7521 (100%)	43,498
–void	5954 (79.2%)	39,541	–retired	1558 (20.7%)	24,821
–void –obsolete	2870 (38.2%)	21,023			
–void –obsolete –incomplete	1855 (24.7%)	8965			

Table 3
AS graph statistics

Source	# of nodes (%inc)	# of edges (%inc)
Oregon route-views	11,174	23,409
+RSs	11,268 (0.84%)	26,324 (12.5%)
+RSs +LG	11,320 (1.3%)	27,899 (19.2%)
+RSs +LG +IRR w/sanity check #1	11,639 (4.2%)	32,903 (40.6%)
+RSs +LG +IRR w/sanity check #2	12,025 (7.6%)	42,639 (82.1%)

dump. In contrast, the second row of Table 3 incorporates *all* available paths. In essence, this row represents the most complete AS graph one can construct from *all* publicly available BGP routing tables. The AS graph reported in the third row was constructed from the AS graph in the second row plus the *Looking Glass* (LG) data. Finally, the AS graph reported in the last two rows include the valid data from the IRR databases sanity-checked by the two different methods described earlier. The “%inc” numbers in parentheses denote the percentage of increase in number of nodes and edges with respect to the Oregon-based AS graph in the first row.

In [21], we reported our other verification efforts confirming that the peering relationships obtained from the IRR reflect *physical* connectivity between ASs.

3.3. AS graph vertex degree distribution revisited

Next, we check how increasingly denser AS graphs affect the power-law characteristics that have been identified by Faloutsos et al. [2] in the Oregon-based AS graphs. Henceforth, we focus on the following AS graphs for comparison: the Oregon-based AS graph (Table 3, first row) and the much denser two AS graphs corresponding to the last two rows in Table 3 (“our AS graphs” or

“our topologies”). We distinguish our two AS graphs by labeling them “our topology I” (the fourth row) and “our topology II” (the fifth row).

We collected nine instances of our data sets, where each instance yields a set of the above-mentioned three AS graphs. These data sets were collected once a week, on the same day of the week, for nine consecutive weeks starting March 2001. With our nine snapshots of the three AS graphs, we plot in Fig. 9 the complementary distribution functions $F^c(x) = 1 - F(x)$, where $F(x)$ is the cumulative distribution function of the AS degree corresponding to one of our nine data sets. “Our” AS graphs I and II are compared against the “Oregon”-based AS graphs in Fig. 9(a) and (b) respectively. In both Fig. 9(a) and (b), all nine instances of “our” AS graphs I and II lie very close to each other and form the upper, curved line in the figure (labeled “Our Topology I” and “Our Topology II”). The nine Oregon-based counterparts also lie very close to each other and form the lower, straight line of the figure (labeled “Oregon Topology”). As is clear from the figures, the vertex degree distributions of the Oregon-based AS graphs appear to be consistent with the strict power-law result reported in [2]. However, the more complete, though not necessarily complete, AS graphs constructed from sources beyond the “Oregon” data set show more ASs with vertex

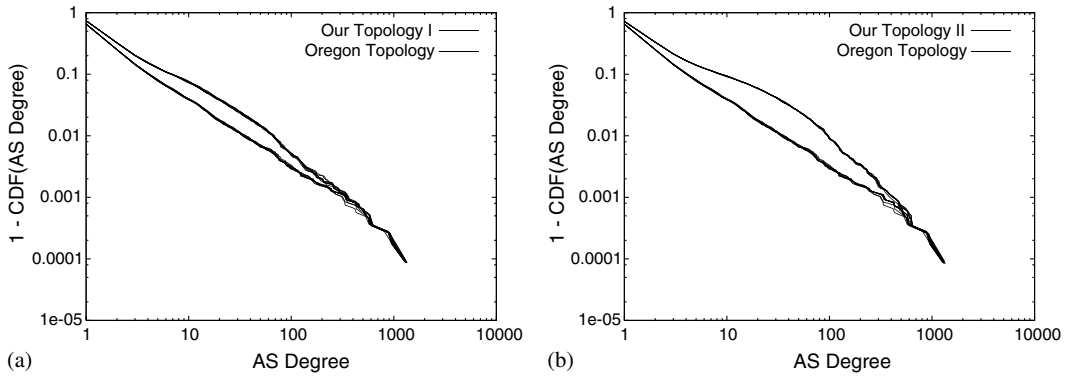


Fig. 9. Frequency distribution of AS degree over time: (a) our topology I and (b) our topology II.

degrees ranging from 4 to 300, resulting in a curved line that deviates from the straight line associated with the Oregon-based AS graphs. Given that two sets of our topologies—obtained from two distinct IRR data sets—share qualitatively the same characteristics, we argue that this observation appears to be insensitive to the specifics of cleaning the IRR data sets.

3.4. The missing links

Fig. 6 in Section 2.3 already established that AS relationships not captured by the Oregon routeviews include both provider-to-customer type and peer-to-peer type relationships. Similarly, we now examine the type of AS peering relationships found on our topology, but not on the Oregon topology. The results based on our topologies I and II agree with the earlier observation. In the following, we present result from our topology I.

Figs. 10 and 11 show that the extra density of our topology is mostly due to peering relationships between ASs located in the mid-level of the AS hierarchy, especially between ASs at the tier-2 and tier-3 levels. The extra density of our topology is due to both provider-customer type (Fig. 10) and peer-to-peer type (Fig. 11) relationships. In these figures, ASs are assigned IDs such that when they are sorted in an increasing order of IDs, they can be grouped by the five-level hierarchy of [22] (i.e., tier-1 to tier-5). For example, s tier-1 ASs are assigned IDs from 1 to s , and t tier-2 ASs are assigned the next consecutive IDs (i.e., $s + 1$ to $s + t$), etc. The boundaries between different hierarchy groups are marked with horizontal and vertical solid lines in the figures. If an AS with ID x_i has a peering relationship with an AS with ID y_j , we place two dots in the appropriate figure, one at coordinates (x_i, y_j) and one at coordinates (y_j, x_i) . Thus the figures are, in essence, graphical repre-

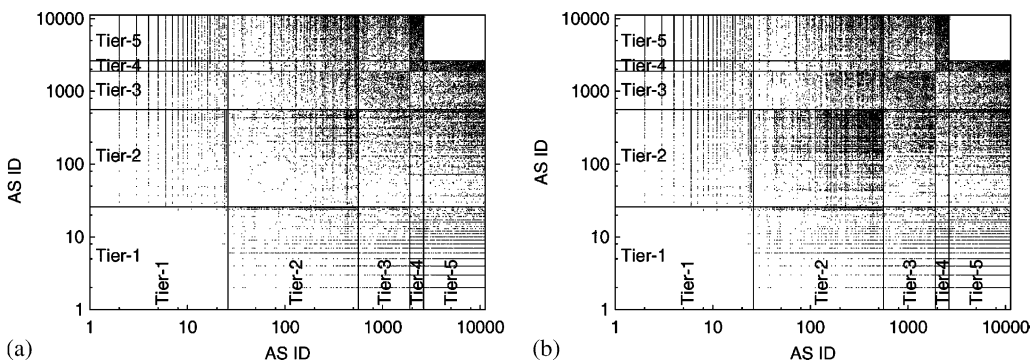


Fig. 10. Hierarchical distribution of provider-to-customer relationships: (a) Oregon topology and (b) our topology.




Fig. 11. Hierarchical distribution of peer-to-peer relationships: (a) Oregon topology and (b) our topology.

representations of AS adjacency matrices, where the density of dots captures the density of peering relationships between ASs. In Figs. 10 and 11, we consider provider–customer type relationships and peer-to-peer type relationships separately. In all, our topology contains up to 180% more peering relationships in the regions (tier- m , tier- n), where $2 \leq m, n \leq 3$, than in the corresponding regions of the Oregon topology; in the other regions, our topology is only about 10% denser than the Oregon topology. This result clearly points to the insufficiency of BGP data, especially for the purpose of inferring peering relationships belonging to mid-tier ASs.

4. Conclusion

The recent increase in research efforts that focus on Internet routing behavior stresses the importance of having access to routing-related measurements. On the one hand, an initiative started by the Oregon route-views project to supply the research community with BGP data has been very instrumental in studying BGP-related phenomena and routing dynamics. On the other hand, due to network security concerns and competitive market conditions, public access to various data sets on network connectivity has been scarce or almost non-existent. As a result, the research community can be expected to experiment with and use routing-related measurements such as the data from Oregon route-views in ways for which the data may not be applicable, sufficient, or intended for.

In this paper, we investigate whether the use of BGP-derived measurements for the purpose of inferring Internet AS connectivity can be justified.

Our results presented in this paper confirm that while actual AS-level connectivity of the Internet is quite high, BGP measurements typically see only a portion of the existing AS connectivity. From BGP's perspective, this observation comes as no surprise. BGP is *not* a mechanism by which ASs distribute their connectivity. It is a protocol by which they distribute the set of routing paths chosen by their policies to reach destinations. Naturally, each AS can only see a subset of those AS connections that are traversed by policy influenced paths.¹³ In this sense, the main lesson learned from the study presented in this paper is that since network-related measurements often reflect network protocol-specific features, arguing for the general validity of an empirical finding about the Internet should be ideally augmented by a careful investigation into the sensitivity of the findings to known deficiencies and inaccuracies of the measurements at hand.

Acknowledgements

We thank Tim Griffin for valuable discussions and for providing us UUNET's BGP routing

¹³ The shortcoming of sampling connectivity at the router-level by collecting shortest path trees has been noted in [37].

table. We also thank the anonymous reviewers for the detailed comments and suggestions that helped improve the presentation of this paper.

References

- [1] University of Oregon Route Views Project, Available from <<http://www.routeviews.org>>.
- [2] M. Faloutsos, P. Faloutsos, C. Faloutsos, On power-law relationships of the Internet topology, in: Proceedings of ACM SIGCOMM, 1999.
- [3] A.-L. Barabási, R. Albert, Emergence of scaling in random networks, *Science* 286 (1999) 509–512.
- [4] R. Albert, A.-L. Barabási, Topology of evolving networks: local events and universality, *Physical Review Letters* 85 (2000) 5234–5237.
- [5] A. Medina, A. Lakhina, I. Matta, J. Byers, BRITE: an approach to universal topology generation, in: Proceedings of MASCOTS, 2001.
- [6] C.R. Palmer, J.G. Steffan, Generating network topologies that obey power laws, in: Proceedings of the Global Internet Symposium, Globecom2000, 2000.
- [7] C. Jin, Q. Chen, S. Jamin, Inet: Internet topology generator, Technical Report UM-CSE-TR-433-00, University of Michigan, EECS Department, Available from <<http://topology.eecs.umich.edu/inet>> 2000.
- [8] H. Tangmunarunkit, R. Govindan, S. Shenker, D. Estrin, The impact of policy on Internet paths, in: Proceedings of IEEE Infocom, 2001.
- [9] H. Tangmunarunkit, R. Govindan, S. Shenker, Internet path inflation due to policy routing, in: Proceedings of SPIE ITCOM, 2001.
- [10] L. Gao, F. Wang, The extent of AS path inflation by routing policies, in: Proceedings of IEEE Global Internet Symposium, 2002.
- [11] K. Park, H. Lee, On the effectiveness of route-based packet filtering for distributed DoS attack prevention in power-law Internets, in: Proceedings of ACM SIGCOMM, 2001.
- [12] Y. Chu, S. Rao, H. Zhang, A case for end system multicast, in: Proceedings of ACM Sigmetrics, 2000.
- [13] National Laboratory for Applied Network Research, NLANR RouteViews archive, Available from <<http://moat.nlanr.net/Routing/rawdata>>.
- [14] Y. Rekhter, T. Li, A Border Gateway Protocol 4 (BGP-4), RFC 1771, March 1995.
- [15] R. Govindan, A. Reddy, An analysis of inter-domain topology and route stability, in: Proceedings of IEEE Infocom, 1997.
- [16] A. Broido, K. Claffy, Internet topology: connectivity of IP graphs, in: Proceedings of SPIE ITCOM, 2001.
- [17] S. Yook, H. Jeong, A.-L. Barabási, Modeling the Internet's large-scale topology, 2001, preprint.
- [18] W. Willinger, R. Govindan, S. Jamin, V. Paxson, S. Shenker, Scaling phenomena in the Internet: critically examining criticality, Proceedings of the National Academy of Sciences 99 (2002) 2573–2580.
- [19] Swiss Network Operators Group, Available from <<http://www.swinog.ch/tools.html>>.
- [20] A. Broido, K. Claffy, Analysis of BGP data from Oregon route views, in: Workshop on Network-Related Data Management, NRDM 2001.
- [21] H. Chang, R. Govindan, S. Jamin, S. Shenker, W. Willinger, Towards capturing representative AS-level Internet topologies, Technical Report UM-CSE-TR-454-02, University of Michigan, EECS Department, 2002.
- [22] L. Subramanian, S. Agarwal, J. Rexford, R. H. Katz, Characterizing the Internet hierarchy from multiple vantage points, in: Proceedings of IEEE Infocom, 2002.
- [23] L. Gao, On inferring autonomous system relationships on the Internet, in: Proceedings of IEEE Global Internet Symposium, 2000.
- [24] E. Chen, J. Stewart, A Framework for Inter-domain Route Aggregation, RFC 2519, February 1999.
- [25] P. Traina, D. McPherson, J. Scudder, Autonomous System Confederations for BGP, RFC 3065, February 2001.
- [26] W. B. Norton, Interconnection strategies for ISPs, unpublished manuscript, 2001.
- [27] Traceroute.org, Public route server and looking glass list, Available from <<http://www.traceroute.org/>>.
- [28] Merit Network, Internet routing registry, Available from <<http://www.irr.net>>.
- [29] C. Alaettinoglu, C. Villamizar, E. Gerich, D. Kessens, D. Meyer, T. Bates, D. Karrenberg, M. Terpstra, Routing Policy Specification Language (RPSL), RFC 2622, June 1999.
- [30] R. Govindan, C. Alaettinoglu, G. Eddy, D. Kessens, S. Kumar, W. Lee, An architecture for stable, analyzable Internet routing, *IEEE Network* 13 (1) (1999) 29–35.
- [31] K. Oberman, J. Haas, Route registry: who uses them? Posting to the NANOG mailing list, October 2000.
- [32] London Internet Exchange (LINX), Memorandum of understanding, Available from <<http://www.linx.net/joining/mou.shtml>>.
- [33] CERN Internet eXchange Point (CIXP), Technical operational environment and principles, Available from <<http://www.wcs.cern.ch/public/services/cixp/cernixp.technical.html>>.
- [34] Internet Neutral Exchange (INEX), Memorandum of understanding, Available from <<http://www.inex.ie/inex-mou.html>>.
- [35] S. McCreary, B. Woodcock, PCH RouteViews archive, Available from <<http://www.pch.net/documents/data/routing-tables>>.
- [36] Merit Network, Internet routing registry mirror site, Available from <<ftp://ftp.radb.net/routing.arbiter/radb/dbase/>>.
- [37] A. Lakhina, J. Byers, M. Crovella, P. Xie, Sampling biases in IP topology measurements, in: Proceedings of IEEE Infocom, 2003.



Hyunseok Chang is a Ph.D. candidate in the Department of Electrical Engineering and Computer Science at the University of Michigan. He received his B.Sc. in Electrical Engineering from Seoul National University, Korea (R.O.K.) in 1998. His current research interests include network measurements, routing, Internet topology, and overlay networks.

Ramesh Govindan received his B.Tech. degree from the Indian Institute of Technology at Madras, and his M.S. and Ph.D. degrees from the University of California at Berkeley. He is an Associate Professor in the Computer Science Department at the University of Southern California. His research interests include Internet routing and topology, and wireless sensor networks.



Sugih Jamin is an Associate Professor in the Department of Electrical Engineering and Computer Science at the University of Michigan. He received his Ph.D. in Computer Science from the University of Southern California, Los Angeles in 1996 for his work on measurement-based admission control algorithms. He spent parts of 1992 and 1993 at the Xerox Palo Alto Research Center, was a Visiting Scholar at the University of Cambridge for part of 2002, and a Visiting Associate Professor at the University of Tokyo for part

of 2003. He received the ACM SIGCOMM Best Student Paper Award in 1995, the National Science Foundation (NSF) CAREER Award in 1998, the Presidential Early Career Award for Scientists and Engineers (PECASE) in 1999, and the Alfred P. Sloan Research Fellowship in 2001.

Scott J. Shenker received his degrees, in theoretical physics, from Brown University (Sc.B.) and the University of Chicago (Ph.D.). After a postdoctoral year at Cornell's physics department in 1983, he joined Xerox's Palo Alto Research Center. He left PARC in 1999 to head up a newly established Internet research group at the International Computer Science Institute (ICSI) in Berkeley. His research over the past 15 years has spanned the range from computer performance modeling and computer networks to game theory and economics. Most of his recent work has focused on the Internet architecture and related issues.



Walter Willinger received the Diploma (Dipl. Math.) from the ETH Zurich, Switzerland, and the M.S. and Ph.D. degrees from the School of ORIE, Cornell University, Ithaca, NY, and is currently a member of the Information and Software Systems Research Center at AT&T Labs—Research, Florham Park, NJ. Before that, he was a Member of Technical Staff at Bellcore (1986–1996). He has been a leader of the work on the self-similar (“fractal”) nature of data network traffic and is co-recipient of the 1996 IEEE W.R.G.

Baker Prize Award from the IEEE Board of Directors and the 1995 W.R. Bennett Prize Paper Award from the IEEE Communications Society for the paper titled “On the Self-Similar Nature of Ethernet Traffic.”