

Towards Capturing Representative AS-Level Internet Topologies

Hyunseok Chang
 Department of EECS
 University of Michigan
 Ann Arbor, MI 48109-2122
hshchang@eecs.umich.edu

Ramesh Govindan
 ICSI
 1947 Center St., Suite 600
 Berkeley, CA 94704-1198
ramesh@icsi.berkeley.edu

Sugih Jamin
 Department of EECS
 University of Michigan
 Ann Arbor, MI 48109-2122
jamin@eecs.umich.edu

Scott J. Shenker
 ACIRI/ICSI
 1947 Center St., Suite 600
 Berkeley, CA 94704-1198
shenker@aciri.org

Walter Willinger
 AT&T Labs-Research
 180 Park Ave.
 Florham Park, NJ 07932-0971
walter@research.att.com

Abstract—Recent studies concerning the Internet connectivity at the AS level have attracted considerable attention. These studies have exclusively relied on the BGP data from Oregon route-views [1] 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 from Oregon route-views may provide only a very sketchy picture of the complete inter-AS connections that exist in the actual Internet has received surprisingly little scrutiny. 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. By augmenting the Oregon route-views data sets 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 connections remain hidden from most BGP routing tables, (2) the AS connections to tier-1 ASs are in general more easily observed than those to non tier-1 ASs, and (3) there are at least about 25–50% more AS connections in the Internet than commonly-used BGP-derived AS maps reveal (but only about 2% more ASs). These findings point out the need for an increased awareness of and a more critical attitude toward the applicability and completeness of given data sets at hand when establishing the generality of any particular observations about the Internet.

Keywords—Internet topology, BGP routing tables

This project is funded in part by NSF grant number ANI-0082287 and by ONR grant number N000140110617. Sugih Jamin is further supported by the NSF CAREER Award ANI-9734145, the Presidential Early Career Award for Scientists and Engineers (PECASE) 1998, and the Alfred P. Sloan Foundation Research Fellowship 2001. Additional funding is provided by AT&T Research, and by equipment grants from Sun Microsystems Inc. and Compaq Corp.

I. INTRODUCTION

For the past two years, there has been a significant increase in research activities related to studying and modeling the Internet’s topology, especially at the level of *autonomous systems* (ASs). For example, these activities include (1) analyzing and modeling measurements to infer the Internet’s AS connectivity graph to describe its properties [2], (2) explaining the origins and causes of some of the observed surprising features [3], [4], (3) building topology generators that produce graph structures that match those of the measured AS connectivity graphs [5], [6], [7], (4) investigating the problem of routing path inflation [8], [9], (5) studying the effectiveness of proposed algorithms for detection/prevention of attacks on (parts of) the network infrastructure [10], and (6) evaluating the performance of multicast protocols [11]. 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 solely for the purpose of collecting their routing tables. The BGP routing tables collected by the Oregon route server are called *Oregon route-views*. From Nov. 1997 to Mar. 2001, the Oregon route-views have been archived on a daily basis by the National Laboratory for Applied Network Research (NLANR) [12]. Presently, archives of the Oregon route-views are available from sites such as the Packet Clearing House (PCH) [13] (starting from Feb. 2001) and *route-views.org* [14] (starting from Apr. 2001).

By making these data sets available to the public, both the Oregon route server and the archival sites are providing invaluable service to the research community. However, the use by researchers of these data sets for the purposes of studying the Internet’s AS connectivity structure

raises the following important issue. The ability to infer the existence of an actual AS connection from BGP routing tables depends largely on existing ASs revealing their relationships since a pairwise business contract may not permit a given AS connection to be available to the third party. In addition, because BGP is a path-vector protocol [15], backup links connecting multi-homed ASs may not show up in BGP routing table snapshots. Consequently, BGP-derived AS connectivity may yield a very incomplete picture of the physical connectivity that exists in the actual Internet. The authors of [16] raise the possibility that BGP-derived AS-level topology snapshots may not be complete and that extracting path information from BGP updates may be a better methodology for obtaining more complete AS topologies. More recently, the router-level connectivity study in [17] suggests that currently available BGP routing tables may not capture many existing AS-level connections. These papers do not attempt, however, to quantify the extent to which the AS topology information derived from BGP snapshots may be incomplete.

Our main objective in this paper is to quantify the completeness of Internet AS maps reconstructed from the Oregon route-views and to attempt to capture *more representative* AS-level Internet topology. 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 AS-level topologies. One of the main contributions of this paper is in developing 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 views (see Section II for a definition of a BGP view and a description of the Looking Glass sites), all originating from different ASs. This BGP-derived connectivity data allow us to explore the question of how well the peering relationships¹ main-

¹Throughout this paper, we will interchangeably use the terms “AS peering relationship,” “AS peering connection” and “AS link.” Two ASs associated with a given AS peering relationship could be connected by a large number of geographically distributed connections at the router level. Additionally, “provider-consumer” relationship or “peer-to-peer” relationship refer to the *contractual* characteristics of a given AS peering relationship. We will also interchangeably use the term “peering AS” and “AS neighbor,” both of which will refer to one of two ASs associated with a given peering relationship with respect to the other.

tained by a given AS (“local view”) are observed by other ASs (“non-local view”). We find that a significant number of existing AS connections, especially those among non tier-1 ASs (see Section II-C for definitions of tier-1 and non tier-1 ASs), are commonly hidden from most BGP routers. We also observe that this phenomenon can be intuitively explained by existing inter-AS peering relationships. In short, the findings reaffirm our earlier comment on the problematic nature of BGP data for the purpose of AS-level topology discovery, and suggest that the actual Internet maintains much richer connectivity at the AS level than has been previously reported.

To quantify the difference between the BGP-derived AS connectivities and the actual inter-AS connections, we consult IRR databases that maintain individual ISP’s routing policy information in several public repositories. The IRR’s goal in maintaining these databases is to coordinate and facilitate the setting of global routing policies. The IRR repository at RIPE contains reasonably up-to-date entries (see Section III-A on how we verify this). We find that AS graphs reconstructed from Oregon route-views, the Looking Glass sites, as well as RIPE information have typically about 25%–50% more edges (and about 2% more nodes) than their counterparts that rely solely on Oregon route-views.

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 must be evaluated in terms of the particular needs of any given study. For example, by themselves, results about routing path inflation, the effectiveness of algorithmic solutions to network security problems, or performance comparisons between different proposed protocols may say little about their 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], that says 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 do suggest a refinement of the original power law observation reported in [2]. More precisely, while our findings state that the vertex degree distributions resulting from more complete snapshots of the AS graph do not conform to the strict power law behavior as stated in [2], they do show that the more complete vertex degree distributions are consistent with the more flexible

class of heavy-tailed distributions that includes the Weibull distribution and the family of distributions where the tail behavior 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 Barabasi 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] (see, however, [18]). Even though the reported constructions can be modified to achieve a better fit to the data and accommodate the observed deviations from a strict power-law behavior (e.g., see [19]), these modifications typically result in highly-parameterized models—a telling sign that the underlying theory provides little physical understanding about the actual Internet topology at the AS level.

The rest of the paper is structured as follows. In Section II, we introduce the notion of a representative BGP view and explore in detail how well peering relationships maintained by an individual AS are being observed by other ASs. To quantify the degree of incompleteness of BGP-derived AS maps, we include in Section III information from the IRR and use its RIPE database to obtain a more complete picture of the physical connections that exist between different ASs (and of which BGP only sees a certain fraction). In Section IV, we seek to provide some intuitive explanations for the observed differences between the BGP-derived AS maps and our more complete picture of the Internet topology at the AS level. We conclude in Section V by commenting on some of the lessons learned and highlighting the implications of our findings.

II. 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, how do we go about checking the completeness of the topology inferred from Oregon route-views? The approach we have adopted in this paper is as follows. The BGP routing table obtained from an AS contains information about that AS’s connectivities to other ASs. It also contains information on the connectivities between other ASs. Assume that the BGP routing table collected at an

TABLE I
BGP DUMP FROM PUBLIC ROUTE SERVERS

Name	AS#	# next hops	# neighbor ASs	Data size (MB)
NC1	-	43	35	291.5
NC2	-	42	16	131.3
C1	7018	24	1	145.8
C2	3967	199	279	127.0
C3	6539	7	1	57.1
C4	3549	3,089	447	49.5
C5	8709	19	187	47.4
C6	1740	3	1	26.0
C7	8220	42	331	20.3
C8	4197	12	82	18.3
C9	3257	1	1	8.0
C10	1	1,175	495	7.8

AS X contains the most complete vertex degree information obtainable of AS X .² The BGP routing table obtained from AS Y will see some, but most likely not all, of the connectivities between AS X and other ASs. Similarly, the BGP routing table obtained from AS Z will see some but not all connectivities 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.

A. Available BGP Routing Tables

In addition to 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 [20]. As of April, 2001, ten commercial ISPs (Internet Service Providers), 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. In addition, 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 de-

²If 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.

note the Oregon and SwiNOG route servers the “collector” route servers, and call the others “operational” route servers. Table I 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 are either the number of iBGP routers connected to a given route server (C1-C3 and C5-C9), or the aggregate number of eBGP routers seen through iBGP routers (C4 and C10). Likewise, the number of “neighbor ASs”. Finally, observe the variability in BGP table sizes (“Data size”). In particular, note the 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.

B. 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 AS-level topology constructed from its BGP 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:

```
* i205.145.52.0 204.255.168.133 701 16758
* i           208.48.18.10    701 3744 16758
*>i          137.39.5.157     701 16758
* i           204.255.168.137 701 16758
```

In this example, the “operational” route server can reach the address space 205.145.52.0 through four of its peer routers, 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 will advertise only the best path, after prepending its own AS 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 our goal had been to infer the AS-level topology from the individual “operational” route servers, 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 III (Table IV).

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 should also be 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

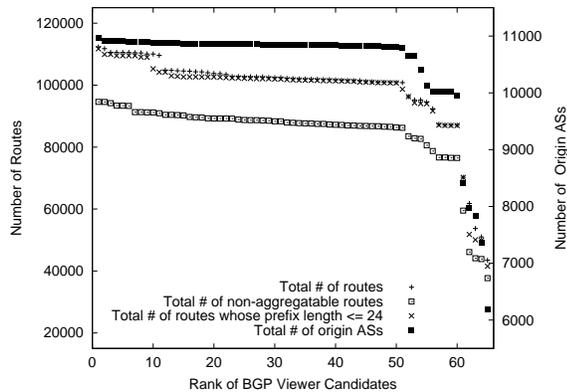


Fig. 1. The Number of Routes (Origin ASs) vs. Rank

number of routes, (2) the number of non-aggregatable routes,³ (3) following [21], 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 server, plus UUNET routing information) in decreasing order and plot the top 65 in Figure 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. These 51 BGP viewers reside in 41 distinct ASs (Table II). Using the BGP views from these 41 ASs in this study gives us 41 perspectives of the Internet. All of the 41 BGP views were collected on the same date (May 25th, 2001) at approximately the same time of day.

C. 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 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. Consequently, for a majority of existing ASs, we essentially have to resort to a rather limited pool of BGP views to predict their connectivities.

To answer the above question, we consider two kinds of BGP views: “local” and “non-local”. From a given AS’s perspective, a BGP view originating from that AS

³An aggregatable route is a redundant route which could be removed from BGP tables by route aggregation; e.g., if a BGP table sees 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.

TABLE II
41 BGP VIEWERS

AS Name (AS#)	Description	Viewer Location
GTE (1)	Backbone	MA, US
STARTAP (10764)	Research network	IL, US
Telstra (1221)	ISP	Australia
VTX (12350)	ISP	Switzerland
SprintLink (1239)	Backbone	VA, US
C&W-Europe (12541)	Backbone	VA, US
INIT7 (13030)	ISP	Switzerland
LAN (15600)	ISP	Switzerland
Telstra USA (16779)	ISP	CA, US
CERFnet (1740)	Backbone	CA, US
Ebone (1755)	Backbone	Sweden
AT&T-GNS (2685)	Backbone	NY, US
XO Comm. (2828)	ISP	CA, US
KPNQwest (286)	Backbone	Europe
ESnet (293)	Research network	CA, US
Tiscali (3257)	ISP	Germany
RIPE NCC (3333)	Internet Registry	Netherlands
C&W (3561)	Backbone	VA, US
Exodus (3967)	ISP	CA, US
Global Online (4197)	ISP	Japan
Globix Corp. (4513)	ISP	OH, US
NETINS Inc. (5056)	ISP	IA, US
JIPPII (5409)	ISP	Germany
Broadwing (6395)	ISP	TX, US
Teleglobe Canada Inc. (6453)	ISP	QC, CA
Abovenet (6461)	ISP	CA, US
SCIFI (6667)	Backbone	Finland
C&W-Switzerland (6893)	Backbone	Switzerland
UUNET (701)	Backbone	NJ, US
AT&T WorldNet (7018)	Backbone	NJ, US
Zocalo (715)	ISP	CA, US
Colt (8220)	Backbone	UK
IBS (8271)	ISP	Switzerland
Teleglobe Europe (8297)	ISP	QC, CA
IXPRIME (8327)	ISP	Switzerland
Dolphins (8758)	ISP	Switzerland
AGRI (8843)	ISP	Switzerland
Carrier1 (8918)	ISP	Europe
ECS-IP (8938)	Backbone	Switzerland
SolNet (9044)	ISP	Switzerland
NEXTRANET (9177)	ISP	Switzerland

is considered *local* and those originating from any other ASs are *non-local*. For example, for AT&T (AS7018), a BGP view from AS7018 is considered local, whereas a BGP view from AS701 is non-local. From the UUNET (AS701)’s perspective, the opposite holds. 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 by its local BGP view. Based on this, we compare—for each of the 41 ASs—AS X ’s vertex degree predicted by its 40 non-local views against that of its local view. By doing so, we will be able to quantify the completeness of *non-local* views.

In Figures 2 to 4, we look at the marginal utility of discovering AS neighbors by 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 degree” in the figures). The opposite order (the smallest AS first) is noted as “increasing AS degree.” The horizontal dotted line in each figure represents the ver-

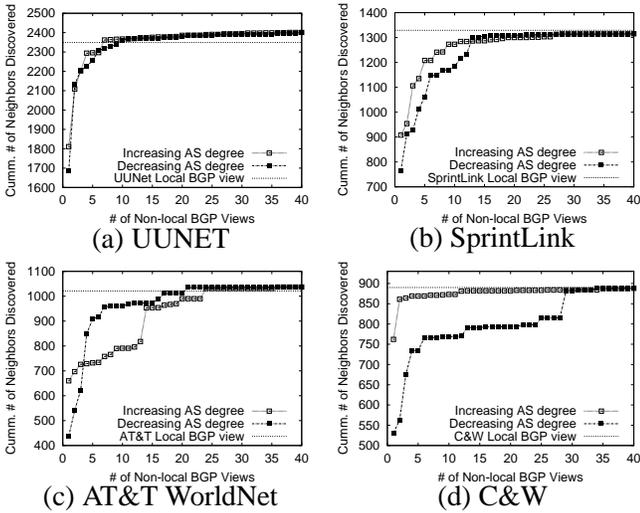


Fig. 2. Marginal Utility of Non-local BGP Views: Tier-1 ASs

tex degree predicted by the given AS’s local BGP view.⁴ Investigating the marginal utility of router-level topology measurements has been recently attempted in [22]. 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. For the purposes of this study, we roughly categorize our 41 ASs into three hierarchy levels (tier-1, tier-2, and tier-3), based on its vertex degree derived from its own local BGP view. The vertex degrees of tier-1 (tier-2) ASs are one order of magnitude larger than those of tier-2 (tier-3) ASs.

Tier-1 ASs. Figure 2 shows that a sufficient number of non-local BGP views can discover most of the neighbors connecting to tier-1 ASs. Interestingly, however, individual non-local views are seen to contribute differently. More specifically, the neighbors of tier-1 ASs tend to be better discovered by non-local views from smaller ASs than from larger ASs. In the case of C&W, the smallest AS’s non-local view discovers 760 out of 890 existing neighbors, whereas the non-local view from the largest AS is only able to discover 530 of them. This phenomenon can be intuitively explained by the *non-transitive* peer-to-peer relationship and the *transitive* provider-consumer relationship [23]. That is, 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

⁴In Figures 2 (a) and (c), 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 AS connections causes some neighbors to be not captured in a given local BGP view snapshot.

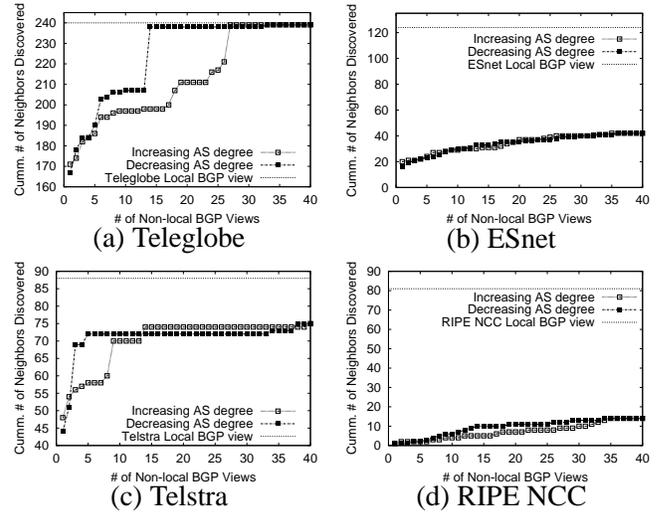


Fig. 3. Marginal Utility of Non-local BGP Views: Tier-2 ASs

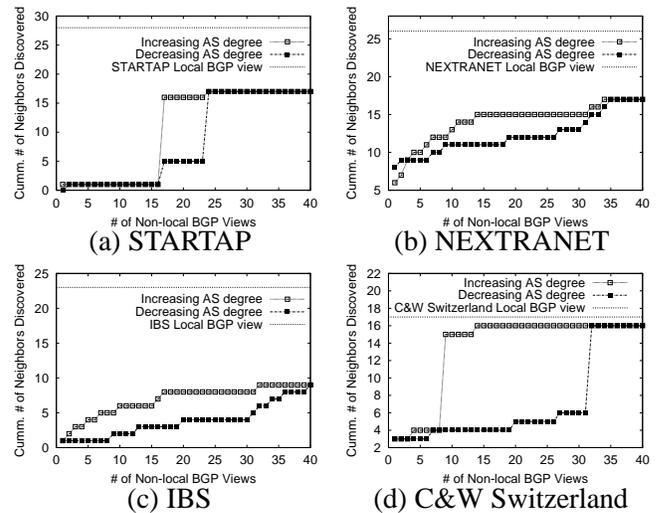


Fig. 4. Marginal Utility of Non-local BGP Views: Tier-3 ASs

downstream customers. For this reason, the BGP views from smaller ASs are likely to better observe tier-1 ASs’ connections.

Non tier-1 ASs. For non tier-1 ASs (Figures 3 and 4), the combined 40 non-local BGP views clearly fail to observe many existing peering relationships, though there are some exceptions (Figures 3(a), 4(a) and 4(d)). In these exceptional cases, a single AS (the 14th largest AS, the 17th smallest AS, and the 9th smallest AS respectively) observes most of the given AS’s neighbors. In the case of Figure 3(a), which is Teleglobe’s AS, the 14th largest AS turns out to be Teleglobe’s own European network. The STARTAP network shown in Figure 4(a) is a National Science Foundation-funded infrastructure connecting several international research sites. The 17th smallest AS in the figure is RIPE NCC, which is one such AS. In the case of Figure 4(d), which is C&W Switzerland, the 9th AS

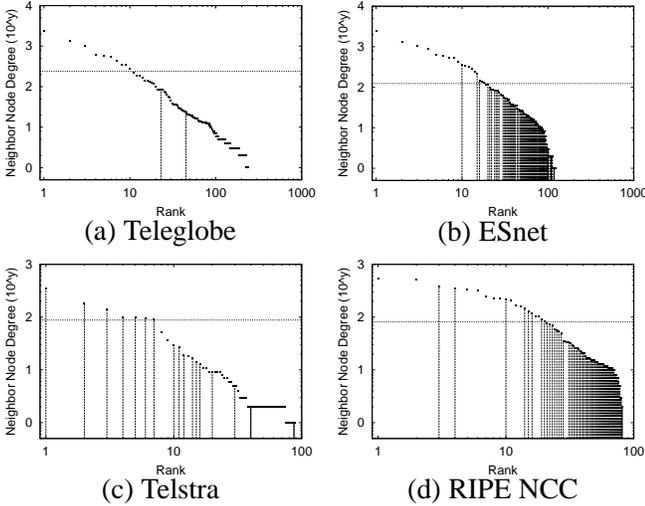


Fig. 5. Vertex Degrees of 4 Tier-2 ASs' Neighbors

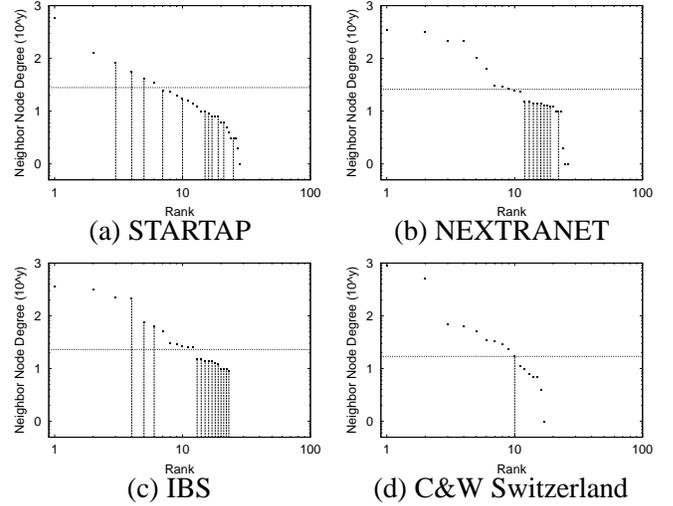


Fig. 7. Vertex Degrees of 4 Tier-3 ASs' Neighbors

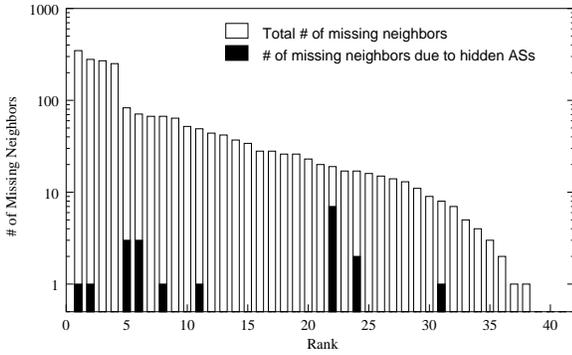


Fig. 6. Number of Missing Neighbors

turns out to be a regional ISP in Switzerland. Conducting a `traceroute` probe to a router in that AS reveals that the ISP has a provider-customer relationship with C&W. These cases re-confirm how well customer views can discover their providers' peering connections. For a majority of non tier-1 ASs whose customer views are not available, their connectivity is not sufficiently approximated by the merged 40 non-local BGP views.

Still, it came as a surprise to us that dozens of BGP views of different ASs are hardly sufficient to capture the major portion of the Internet connectivity as far as non tier-1 ASs are concerned. Given that a non-negligible number of neighbors of a given AS 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 to other ASs as part of a larger aggregated address space [24]. On the other hand, a hidden *edge* means that while the AS numbers of both

end points of the edge are present in other ASs' routing tables, the neighbor peering relationship between the two end points is not listed in any AS path. Figure 6 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 caused by the non-local views not detecting the peering relationships between AS X and its neighbors. One can see that the number of missing neighbors from hidden ASs is negligibly small (the y -axis is in log-scale). Thus, the majority of missing neighbors are caused by hidden AS *links*, not hidden ASs.

Given that BGP views can fail to observe a non-negligible number of existing AS links, we want to know what differentiates those links from the visible ones. In the absence of detailed AS relationship and internal policy routing information, we base our analysis on studying AS vertex degree as follows. Since most missing neighbors occur with respect to non tier-1 ASs, we only consider non tier-1 ASs. Figures 5 and 7 show the vertex degrees of several non tier-1 ASs' neighbors, sorted in decreasing rank order.⁵ The horizontal line in the figures indicate the vertex degree of the local AS (for example, Figure 5(a) shows that Teleglobe has a vertex degree of about 250). A neighbor that is hidden in all non-local views is marked with a vertical dashed line. The frequency of vertical dashed lines indicates how many neighbors are missing in non-local views; the y values of those lines record the vertex degrees of the missing neighbors. Except for Figure 5(c), the first few largest neighbors (most likely upstream providers) of each AS are well observed by other ASs. This is because AS links connecting an AS and

⁵An AS vertex degree here is derived from the merged 41 BGP views.

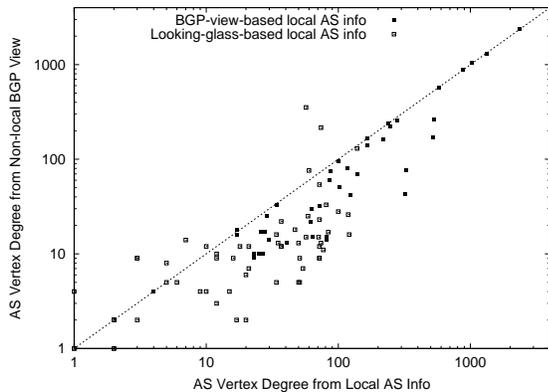


Fig. 8. The Completeness of Non-local BGP Views

its upstream providers should eventually be visible in the BGP views of tier-1 ASs that are located in the top level of the AS hierarchy. Thus the majority of missing neighbors have vertex degrees comparable to or less than that of the given AS. We conjecture that many of these hidden links are private peer-to-peer relationships. See Section IV for observations supporting this conjecture.

D. Other BGP-derived Connectivity Information

To help troubleshoot Internet-wide routing problems, several ISPs make available 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 aggregated BGP statistics for each of them. From this BGP summary information, we can elicit the set of AS neighbors connected to the local AS.⁶

Once we obtain the number of each AS’s neighbors from its corresponding BGP summary information, we compare it with the one predicted by our BGP views, as in Section II-C. However, unlike the BGP routing tables of our BGP viewers, the BGP summary information from individual Looking Glass sites may not list all existing neighbors of corresponding ASs, since those routers originating the information may cover only limited address space.⁷ Therefore, the number of neighbors revealed from an individual piece of BGP summary information can only be interpreted as a lower bound on the number of existing neighbors connecting to a given AS.

⁶Querying a Looking Glass is done through web-based interface; we pre-selected some sites [25] and have our crawling script periodically collect their BGP summary information. We started our script at the same time we collected BGP dumps from route servers.

⁷When a BGP router does not have an AS path to a given address space, it can forward all packets addressed to that address space to a default router.

Figure 8 visualizes 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. Two labels in the figure identify which source is used to obtain the actual degree of a given AS. The black dots labeled “BGP-view-based local AS info” summarize the neighbor discovery results of Section II-C. 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 Figure 8 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 fall short.

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

III. 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 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 resort to the Internet Routing Registry (IRR) [26] 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 at various levels (e.g., individual address prefixes or ASs, etc.). The following two hypothetical database records show how such routing information is expressed using the Routing Policy Specification Language (RPSL).

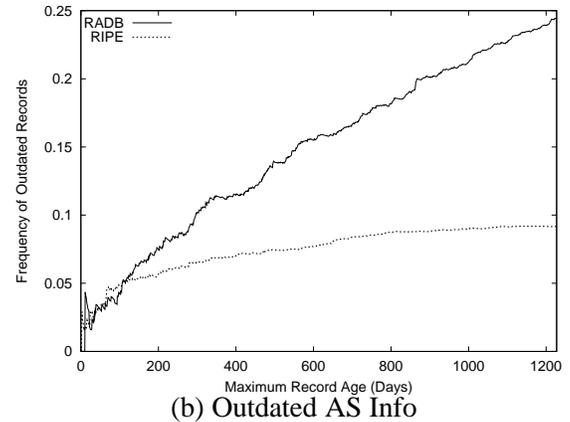
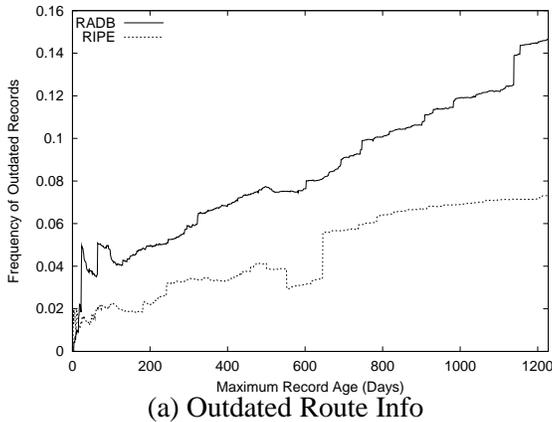


Fig. 9. Frequency of Outdated Records

```

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 record states that “1.2.3.0/24” belongs to AS1 as of Mar. 13, 2001. The latter record, 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. From this import and export policies specification, we can infer the neighboring ASs of AS1.

A. On the Freshness of the Internet Routing Registry

We next question the reliability of such manually-registered policy routing information. 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 [27]. However, being predicated on 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 [28], 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 [29], [30], [31] specify as a membership requirement that

members register their routes and peering policy in the RIPE database. Given this, we consider the RIPE database a potentially reliable source of AS-level connectivity. For example, when comparing the RADB and RIPE databases of May 25, 2001, we found that while out of the 2,673 ASs registered with RADB, only 2,039 (76.3%) published their routing policy, 4,203 (93.6%) out of the 4,492 ASs that had registered with RIPE published their routing policy.

To further verify the relative freshness of the RIPE database, we checked the individual records of the RADB and the RIPE database as follows. For each routing registry record, we looked at its last update time and correlated its routing information with the Oregon route-views [12], [13] corresponding to the *same* period as the update time. Registry entries recording specific route (address prefix) were checked for their route origin information; registry entries specifying the routing policies of ASs were checked against the ASs’ neighbors lists.

We downloaded public IRR database files mirrored at [32] on May 25th 2001 and compared each of their records with the Oregon route-views collected since the record’s last update time. A given registry entry can be either (a) consistent, (b) inconsistent, or (c) not available in the tables. For example, a BGP table can correctly observe that the prefix 1.2.3.0/24 originates from AS X (a), or incorrectly indicates that 1.2.3.0/24 belongs to AS Z (b), or simply does not observe the prefix at all (c). We consider a registry entry outdated if its information used to be consistent with an older BGP routing table⁸ but has become inconsistent with or not available in the more recent (May, 25th) BGP table.

⁸A record is considered consistent with a BGP table if the route origin AS specified in the record is correct according to the BGP table, or if a given AS described by the record has registered all its peering relationships found in the BGP table.

Figure 9 compares RADB and RIPE in terms of freshness. They show the frequency (y -axis) of outdated records 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 a record thus indicates how many days have passed between the time the record was last updated and May 25th. The figure clearly demonstrates that the RIPE database is maintained more carefully and in a more up-to-date manner than the RADB database.

B. Obtaining AS Connectivity from Registry Data

Based on the relative completeness and freshness of the RIPE database, we decided to use it in our study. In order to avoid including any incorrect or outdated information from the RIPE database, we applied the following more stringent checks to individual database records:

Void records. We consider an IRR record *void* if the AS described by the record was once present in the Oregon route-views (dating from Nov. 1997) but has disappeared from the tables afterwards.⁹

Obsolete records. To find *obsolete* records, we first construct an *AS reference graph* from the registry records. 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 one neighbor. An edge is directed from node A to node B if the registry record of AS A specifies AS B as a peering neighbor. If all the records were up-to-date, all edges on the 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 record updated earlier *obsolete*.

Incomplete records. We consider a record *incomplete* if the AS described has a neighbor in the Oregon route-views that was not registered with IRR. To detect unregistered neighbors, we scanned the daily Oregon route-views starting from the date of the record’s last update to the end of our data set (May 25, 2001). If an unregistered neighbor of an AS is found, the AS’s record in the IRR is considered incomplete.

Only records that are not void, not obsolete, and not incomplete are considered *valid* records. Table III shows that only 1,026 ASs (about 9% of all known ASs) have *valid* records in the RIPE database. Nevertheless, as Table IV shows, perusing the IRR database allows us to identify an

⁹To 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 (May 25, 2001).

TABLE III
VALIDITY ANALYSIS OF RIPE DATABASE

	# of records (ASs)
All records	4,203
– void	3,917
– void – obsolete	1,582
– void – obsolete – incomplete	1,026

TABLE IV
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 + RIPE	11,456 (2.5%)	32,759 (40.0%)

extra 4,860 edges (or about 17.42%) over the most complete AS map constructed from all the BGP information we can obtain (compare the last and second-to-last rows of Table IV). The table shows the number of nodes (ASs) and edges contained in the AS map constructed from the various sources, cumulatively. The first row, labeled “Oregon route-views” lists the number of nodes and edges found in the AS map constructed from 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 I. Recall that in the BGP view analysis in Section II, we used only the best paths from each full BGP dump. In contrast, the second row of Table IV incorporates *all* available paths. In essence, this row represents the most complete AS map one can construct from *all* publicly available BGP routing tables. The AS map reported in the third row was constructed from the AS map in the second row plus the *Looking Glass* (LG) data. Finally, the AS map reported in the last row includes the valid data from the RIPE database. The “%inc” numbers in parentheses denote the percentage of increase in number of nodes and edges with respect to the Oregon-based AS map of the first row.

C. AS Graph Vertex Degree Distribution Revisited

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

We collected 9 instances of our data sets, where each instance yields a pair of the above-mentioned two AS graphs. These data sets were collected once a week, on

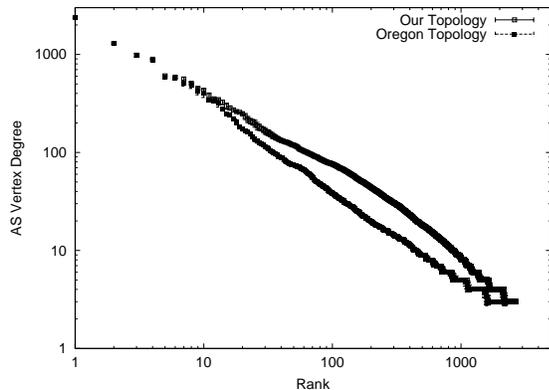


Fig. 10. Power-Law 1: AS Vertex Degree vs. Rank

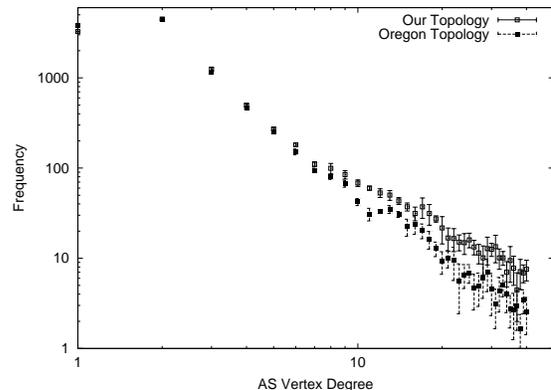


Fig. 11. Power-Law 2: Frequency Distribution of AS Vertex Degree

the same day of the week, for 9 consecutive weeks starting Mar. 2001. With our 9 snapshots of the two AS graphs, we plot in Figure 10 and Figure 11 the time-averaged vertex-degree-rank distribution and vertex-degree-frequency distribution respectively. In the figures, we also plot the standard deviation from the mean.

According to Figure 10, the newly-added edges found in our topology significantly increase the vertex degree of nodes with ranks between 10 and 1000. Thus, the overall vertex-degree-rank distribution takes on a more pronounced nonlinear (i.e., convex) shape. The frequency distribution of Figure 11 shows that newly-added edges have minimal effect on the frequency of nodes whose vertex degrees are less than 10. However, starting from vertex degree 10 or so, the frequency distributions resulting from the two types of AS graphs start to significantly deviate from one another—the confidence intervals of the form $[\text{mean} - \sigma, \text{mean} + \sigma]$ (here, σ denotes the standard deviation resulting from each of the 9 snapshots¹⁰) associated with degrees 10 and larger do not overlap! These results are consistent with our earlier findings in II-C suggesting that there may exist much richer connectivity among non tier-1 ASs than observed by a handful of BGP routers.

In Figure 12 we plot 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 9 data sets. All 9 instances of our AS graph lie very close to each other and form the upper, curved line in the figure (labeled “Our Topology”). The 9 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 figure, the vertex degree distributions of the Oregon-based AS graphs appear to be

¹⁰While datasets collected over a longer than 9-week period would yield better estimates of σ , such datasets would also start to be too far apart in time, capturing therefore possibly quite different configurations of the continuously growing Internet.

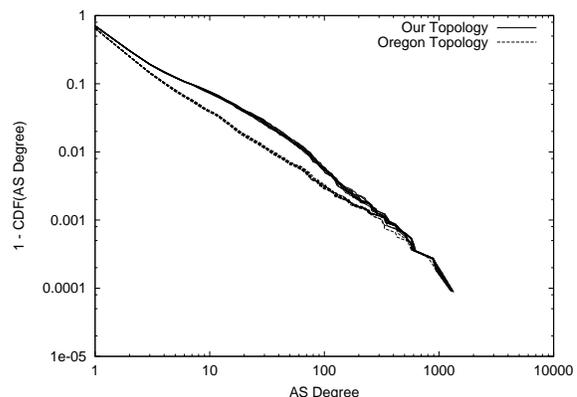


Fig. 12. Frequency Distribution of AS Degree Over Time.

consistent with the strict power-law result reported in [2]. However, the more complete, though not necessarily complete, AS maps constructed from sources beyond the “Oregon” data set show more ASs with vertex degrees ranging from 4 to 300, resulting in a curved line that deviates significantly from the straight line behavior associated with the Oregon-based AS graphs. However, the vertex degree distributions of “our” AS graphs are clearly heavy-tailed or highly-variable in the sense that the observed vertex degrees typically range over three or four orders of magnitude.¹¹

IV. ON VALIDATING AND GENERALIZING OUR AS GRAPH

In this section, we address two questions regarding our much denser AS graph. First, could our AS graph be artificially inflated by our use of non BGP-derived con-

¹¹Performing a more detailed analysis (not shown here), we found that while the degree distributions resulting from the Oregon-based AS graphs are consistent with strict power-law distributions, the corresponding distributions of our more complete AS graphs are not. More precisely, the degree distributions of our AS graphs are consistent with heavy-tailed distributions such as the Weibull distribution [17] or distributions where only the tail obeys a power-law [18].

TABLE V
AS CO-LOCATION INFORMATION OF EUROPEAN EPS

EP	# Co-located ASs	Location
1	129	London, UK
2	114	Amsterdam, The Netherlands
3	73	Vienna, Austria
4	53	Paris, France
5	49	Germany
6	44	Zurich, Switzerland
7	44	Brussel, Belgium
8	39	Oslo, Norway
9	38	Budapest, Hungary
10	38	Milan, Italy
11	32	Copenhagen, Denmark
12	30	Slovakia
13	28	Madrid, Spain
14	20	Stockholm, Sweden
15	19	Zurich, Switzerland
16	13	Helsinki, Finland

nectivity information such as IRR-based data, especially since the latter may not reflect actual physical connectivity? Second, since we use only the European RIPE registry database, we cannot claim that our AS graph is representative of the *global* AS-level Internet topology. Would our observations on the AS vertex-degree distribution in Section III-C still be valid on the actual AS graph?

A. Validating AS-level Connectivity

According to Table IV, our AS graph has 9,350 (40%) more edges than the Oregon-based counterpart. Checking whether these newly found edges represent actual physical connectivity requires identifying whether ASs that are connected by any of these newly discovered edges are in fact physically connected to each other. Our test for this is based on identifying those ASs that are *physically co-located* at existing public exchange points (EPs). There are about 70 EPs in Europe, of which 16 publish their co-location information through Looking Glass. Table V lists these 16 EPs and the number of ASs co-located at each of them. Among the 9,350 new edges, 4,811 of them (51.5%) occur at one of these 16 EPs. This percentage should be considered significant given that we investigate only 16 of the 70 or so EPs in Europe. This high percentage not only supports the validity of newly found *individual* edges in our AS graph, but also suggests that the relative density of our graph is a result of our capturing connectivities at EPs.

If it were true that most of the connections observed in our AS graph, but not in the Oregon topology, are due to ASs' connecting to each other at EPs, we should expect to see large cliques in our graph. The presence of large cliques might indicate the existence of connectivity hubs (EPs). We next look for cliques in our AS graph, study their sizes relative to those of the Oregon topology, and determine whether they occur at existing EPs.

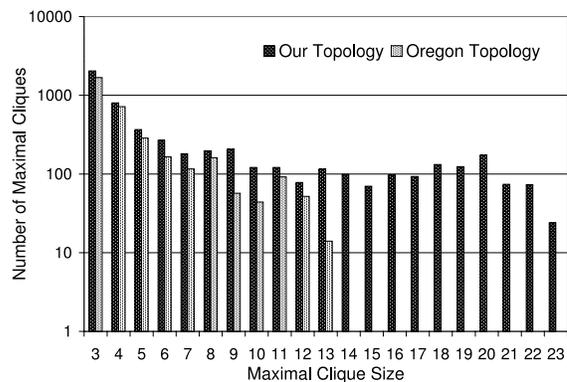


Fig. 13. Maximal Clique Distribution

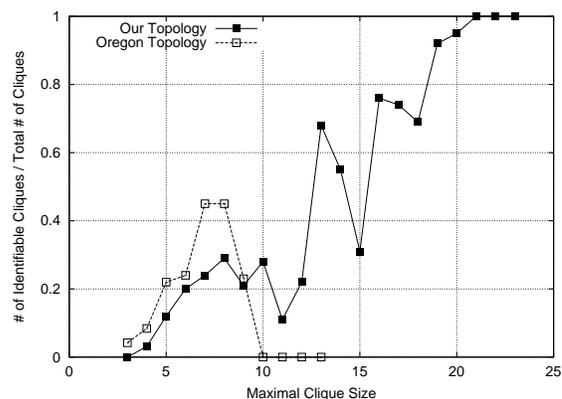


Fig. 14. Identifiable Maximal Cliques

A *maximal* clique is a clique which is not contained in any other clique. A *maximum* clique is a maximal clique of the largest size. Our clique search algorithm, which searches for maximal cliques, incrementally grows an initial clique by degree-based greedy heuristic. The initial clique is any single node whose degree is greater than or equal to three. We run our search algorithm on both our AS graph and its Oregon-based counterpart (labeled as “Our Topology” and “Oregon Topology” in Figure 13). According to the figure, the largest size of maximal cliques found in Oregon-based AS graph and our AS graph is 13 and 23 respectively. Given the much richer connectivity of our AS graph, finding much larger cliques in our AS graph is not surprising.

Now we examine whether these individual maximal cliques consist of co-located ASs. If we can find an EP where *all* ASs in a given maximal clique are co-located, we consider the clique *identifiable*. Figure 14 shows the ratio of identifiable cliques over all the maximal cliques found in Figure 13. Once again, we consider only the 16 European EPs listed in Table V as candidate exchange points. We observe that many large cliques captured in our AS graph are identifiable, and hence can be justified (notably,

all of those with sizes 21 to 23).¹² This observation also re-confirms that Oregon-based AS graphs fail to include a significant portion of existing AS-level connectivities.

The validation results of our AS graph so far strongly suggests that the added connectivity of our AS graph reflects the actual existing AS-level connectivity.

B. Generalizing Our Findings

We now address the question of whether our observations on the AS vertex-degree distribution in Section III-C can be expected to hold for more global versions of our AS graphs.

These are the findings we have made so far: (1) the total number of currently existing ASs is already well predicted by both our AS graph and the Oregon-based version (Table IV). (2) the complete connectivity of tier-1 ASs is contained in our AS graph (Section II). (3) the connectivity of end-customer ASs, which would typically be a small number of provider-customer connections, can be sufficiently approximated by both our AS graph and Oregon-based counterpart. (4) *our use of the European RIPE database enables capturing the dense connectivity occurring at European EPs.*

Figure 15 shows how many ASs are co-located at each EP in different regions of the world.¹³ While the numbers of co-located ASs at the largest EPs in North America are lower than those in the rest of the world, we typically see some 10–100 ASs co-located at each of the different EPs. Hence we believe that a complete map of the *whole* (i.e., global) AS-level topology will result in a denser graph across the world.

Referring back to Figure 12, the implication of these findings are as follows: the first and third observations prevent the frequencies of the smallest degrees (the head of the degree frequency distribution curve, Figure 12) from increasing significantly. The second observation pins down the tail of degree frequency distribution curve as is. Considering that all of the EPs, not just the ones in Europe, have 10 to 100 ASs co-located at each of them, the fourth observation implies that we will have a *larger number of ASs with denser connectivity* than what is captured by Oregon route-views. As a result, the middle portion of the degree frequency distribution curve (Figure 12) can be expected to experience a further shift upwards and to

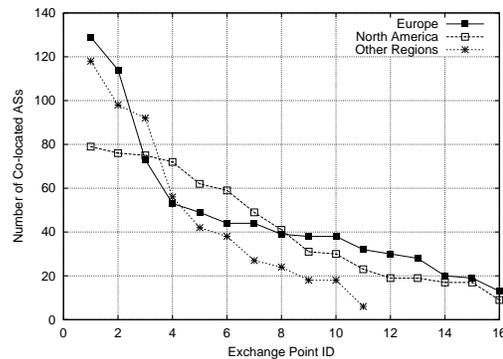


Fig. 15. Number of Co-located ASs at EPs Worldwide. As of Sept, 2001, there are at least 68 EPs in Europe, 40 in Asia/Pacific, 58 in North America, 14 in Latin America, and 7 in Africa/Middle East region [33].

the right. A combination of these qualitative changes can therefore be expected to cause the degree frequency distribution to become more nonlinear (i.e., curving further *outward*), with its head and tail remaining essentially unchanged. Clearly, the resulting distribution will remain *heavy-tailed* or *highly-variable*, but will very likely no longer conform to the strict power-law behavior that characterizes the Oregon-based AS maps.

V. CONCLUSION

With the recent significant increase in research efforts focusing on Internet routing behavior, routing-related measurements have become a highly valuable commodity, and NLANR and the Oregon route server have been the leaders in supplying the research community with invaluable and relevant data, in particular with BGP measurements. Compared to other concurrent efforts that have focused, for example, on Internet workload characterization, the number of researchers with expertise in Internet routing is still relatively small. As a result, the general research community can be expected to be more prone to using routing-related measurements such as the Oregon route-views in ways for which the data are really not applicable, sufficient, or intended for, in the first place.

In this paper, we show that indeed, as far as past investigations into the Internet topology at the AS level are concerned, many of the findings that have been reported in the literature have used the publicly available BGP measurements without realizing the possible pitfalls associated with taking the data at face value, or without examining whether or not the use of the data is justified for inferring the Internet AS connectivity.

Our results confirm that while the actual connectivity of the Internet at the AS level is quite high, BGP measurements typically see only a portion of all existing AS

¹²The reason that most of the large-sized cliques (10 to 13 ASs) of Oregon-based graph are not identifiable is that the ASs involved are mostly tier-1 ASs, which are not co-located at the 16 European EPs.

¹³For non-European EPs, we collected the number of co-located ASs either from Looking Glass queries or published documents on the EPs' web sites. When *validating* our AS graph, however, we do not use published co-location information.

connections. From BGP's perspective, this observation comes as no surprise, because it expresses the defining property of BGP; that is, BGP reflects AS relationships and not physical AS connectivity. 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 typically include a careful investigation into the sensitivity of the findings to known deficiencies and inaccuracies of the measurements at hand.

REFERENCES

- [1] "University of Oregon Route Views Project," <http://www.anc.uoregon.edu/route-views/>.
- [2] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On power-law relationships of the Internet topology," in *Proceedings of ACM SIGCOMM '99*, August 1999.
- [3] A.-L. Barabasi and R. Albert, "Emergence of scaling in random networks," *Science*, vol. 286, pp. 509–512, Oct 1999.
- [4] R. Albert and A.-L. Barabasi, "Topology of evolving networks: Local events and universality," *Physical Review Letters*, vol. 85, pp. 5234–5237, 2000.
- [5] A. Medina, A. Lakhina, I. Matta, and J. Byers, "BRIT: An Approach to Universal Topology Generation," in *Proceedings of MASCOTS '01*, August 2001.
- [6] C. R. Palmer and J. G. Steffan, "Generating network topologies that obey power laws," in *Proceedings of the Global Internet Symposium, Globecom2000*, November 2000.
- [7] C. Jin, Q. Chen, and S. Jamin, "Inet: Internet topology generator," Tech. Rep. CSE-TR-433-00, University of Michigan, EECS Dept., 2000, <http://topology.eecs.umich.edu/inet>.
- [8] H. Tangmunarunkit, R. Govindan, S. Shenker, and D. Estrin, "The Impact of Policy on Internet Paths," in *Proceedings of IEEE Infocom*, 2001.
- [9] H. Tangmunarunkit, R. Govindan, and S. Shenker, "Internet Path Inflation Due to Policy Routing," in *Proceedings of SPIE ITCOM 2001*, August 2001.
- [10] K. Park and H. Lee, "On the effectiveness of route-based packet filtering for distributed DoS attack prevention in power-law Internets," in *Proceedings of ACM SIGCOMM '01*, 2001.
- [11] Y. Chu, S. Rao, and H. Zhang, "A case for end system multicast," in *Proceedings of ACM Sigmetrics*, 2000.
- [12] National Laboratory for Applied Network Research, "NLANR RouteViews archive," <http://moat.nlanr.net/Routing/rawdata>.
- [13] S. McCreary and B. Woodcock, "PCH RouteViews archive," <http://www.pch.net/documents/data/routing-tables>.
- [14] Routeviews.org, "Route Views archive," <http://archive.routeviews.org>.
- [15] Y. Rekhter and T. Li, "RFC 1771: A Border Gateway Protocol 4 (BGP-4)," Mar 1995.
- [16] R. Govindan and A. Reddy, "An Analysis of Inter-Domain Topology and Route Stability," in *Proceedings of IEEE Infocom*, 1997.
- [17] A. Broido and k. claffy, "Internet topology: connectivity of IP graphs," in *Proceedings of SPIE ITCOM*, August 2001.
- [18] W. Willinger, R. Govindan, S. Jamin, V. Paxson, and S. Shenker, "Scaling phenomena in the Internet: Critically examining criticality," in *Proceedings of the National Academy of Sciences*, 2001.
- [19] S. Yook, H. Jeong, and A.-L. Barabasi, "Modeling the Internet's large-scale topology," 2001, preprint.
- [20] "Swiss Network Operators Group," <http://www.swinog.ch/tools.html>.
- [21] A. Broido and k. claffy, "Analysis of BGP data from Oregon route views," *Workshop on Network-Related Data Management (NRDM 2001)*, May 2001.
- [22] P. Barford, A. Bestavros, J. Byers, and M. Crovella, "On the Marginal Utility of Network Topology Measurements," November 2001, To appear in the SIGCOMM Internet Measurement Workshop.
- [23] L. Gao, "On inferring autonomous system relationships on the Internet," in *Proc. IEEE Global Internet Symposium*, November 2000.
- [24] E. Chen and J. Stewart, "RFC 2519: A framework for inter-domain route aggregation," February 1999.
- [25] Traceroute.org, "Public route server and looking glass list," <http://www.traceroute.org/>.
- [26] Merit Network, "Internet routing registry," <http://www.irr.net>.
- [27] R. Govindan, C. Alaettinoglu, G. Eddy, D. Kessens, S. Kumar, and W. Lee, "An architecture for stable, analyzable Internet routing," *IEEE Network Magazine*, January-February 1999.
- [28] K. Oberman and J. Haas, "Route registry: who uses them?," October 2000, Posting to the NANOG mailing list.
- [29] The London Internet Exchange (LINX), "Memorandum of understanding," <http://www.linx.net/joining/mou.html>.
- [30] The CERN Internet eXchange Point (CIXP), "Technical operational environment and principles," <http://www.wcs.cern.ch/public/services/cixp/cernixp.technical.html>.
- [31] The Internet Neutral Exchange (INEX), "Memorandum of understanding," <http://www.inex.ie/inexmou.html>.
- [32] Merit Network, "Internet routing registry mirror site," <ftp://ftp.radb.net/routing.arbiter/radb/dbase/>.
- [33] EP.NET, "A global list of public exchange points," <http://www.ep.net>.