

Collecting the Internet AS-level Topology

Beichuan Zhang, Raymond Liu
Computer Science Dept.
UCLA
{bzhang, raymondli}@cs.ucla.edu

Daniel Massey
Computer Science Dept.
Colorado State University
massey@cs.colostate.edu

Lixia Zhang
Computer Science Dept.
UCLA
lixia@cs.ucla.edu

ABSTRACT

At the inter-domain level, the Internet topology can be represented by a graph with Autonomous Systems (ASes) as nodes and AS peerings as links. This AS-level topology graph has been widely used in a variety of research efforts. Conventionally this topology graph is derived from routing tables collected by RouteViews or RIPE RIS. In this work, we assemble the most complete AS-level topology by extending the conventional method along two dimensions. First, in addition to using data from RouteViews and RIPE RIS, we also collect data from many other sources, including route servers, looking glasses, and routing registries. Second, in addition to using routing tables, we also accumulate topological information from routing updates over time. The resulting topology graph on a recent day contains 44% more links and 3% more nodes than that from using RouteViews routing tables alone. Our data collection and topology generation process have been automated, and we publish the latest topology on the web on a daily basis.

1. INTRODUCTION

The Internet consists of tens of thousands of loosely connected networks called Autonomous Systems (ASes) and the Border Gateway Protocol (BGP [15]) is used to exchange reachability information among Autonomous Systems. The entire Internet can be viewed as an *AS-level topology graph* where each AS is a node, and the BGP peering between two ASes is a link¹. This AS-level topology has important implications for both day to day Internet operations and Internet research. Estimates of the AS-level topology have been used in a variety of research activities, including analyzing Internet topological properties [8][7], inferring AS relationship and Internet hierarchy [9][18][5], building network topology generators [13] for simulations, and evaluating the effectiveness of new protocols and improvements [14].

¹In this paper, we treat each link in the topology as a logical peering relationship between two ASes, instead of a physical link between two routers.

An AS-level topology estimate can be derived from BGP routing tables. Each entry in a BGP routing table lists the path of ASes used to reach a destination prefix, and thus each entry implicitly lists AS connectivity information. An estimate of the Internet AS-level topology can be obtained by taking the union of all AS paths found in BGP routing tables available from either RouteViews [17] or RIPE RIS [16]. However, each router can only see the Internet connectivity from its own limited view. By including additional data from route servers, looking glasses, and routing registry databases, [6] showed that the topology estimate can be significantly augmented. Furthermore, a BGP routing table snapshot only captures AS links used in the router's BGP paths at the time, even though multiple alternative paths exist for almost all the destinations. As we will show in this paper, a topology estimate can be significantly improved by routing updates, which over time expose alternative paths when primary paths become unavailable.

However collecting a topology from many sources and over time requires a significant amount of time and work. Researchers often settle for a topology estimate obtained using a single snapshot of RouteViews routing tables. A better solution would be to produce a freely available AS-level topology estimate that incorporates multiple data sources and multiple time periods.

In this work, our goal is to capture the AS-level topology to the furthest extent using existing resources and make this topology widely available. More precisely, the topology should have the following properties:

- **Most Complete:** The topology information should be collected from as many inter-domain data sources as possible using state-of-the-art methodologies. Currently, our data sources include RouteViews, RIPE, route servers, looking glasses, and routing registries, and we accumulate topological information from routing updates over time.
- **Annotated:** Auxiliary information should be included to help better use of the topology. For example, a fundamental trade-off in collecting AS-level topology is “completeness” vs. “freshness.” Using routing registry data and accumulating information over time make the topology more complete, but may also introduce outdated information. Associating each node or link with a timestamp will give users the flexibility to decide how

	timestamp	type	peer IP	peer AS#	prefix	AS path	others
BGP4MP	1067645344	B	10.0.0.1	123	131.179.0.0/16	123 456 789 987	...

Table 1: A BGP Routing Table Entry in Trace Data

“fresh” they want the topology be. Other auxiliary information includes the type of nodes and links, and from which data sources they are collected.

- **Up-to-date:** Since the Internet changes every day, the topology collection process should be automated and update the topology on a daily basis.
- **Easily Available:** The topology should be published on the web and include both the current estimate and past topology estimates.

We use multiple inter-domain level data sources and accumulate topological information over time to provide the *most complete, annotated, up-to-date*, and *easily available* AS-level topology. We will present the resulting topology on October 24, 2004, which contains 44% more links and 3% more nodes than the one solely derived from RouteViews routing tables. Our topology is updated daily and available on the web (<http://irl.cs.ucla.edu/topology/>). In addition to the most recent topology, topologies from previous days are also available to provide a historical record and enable longitudinal studies.

We believe that making the most complete and up-to-date topology easily available will benefit the research community as a whole. However researchers should also keep in mind the inherent limitations in collecting network topologies. The real Internet topology is unknown and *any* AS-level topology estimate, including the one presented in this paper, is not necessarily complete. Adding even more data sources and looking over longer time periods may result in diminishing return in topology collection, but this does not imply the topology is near complete. For example, some peer to peer AS links may only be revealed by monitors in particular locations or may only be revealed when exceptional network events occur. Research results should consider topology (in)completeness, nevertheless a significantly more complete topology may benefit many research efforts.

The rest of the paper is structured as follows. Section 2 describes our dataset, its sources and how to extract topological information from raw data. Section 3 discusses how to compile a topology from different pieces of information, and compare the contributions of different data sources. Section 4 presents the final topology on October 24, 2004, and its degree distributions of different types of nodes. Section 5 discusses the related work, and Section 6 concludes the paper.

2. DATA SOURCES

The sources that we collect raw data from fall into four categories: BGP trace collectors, route servers, looking glasses, and the Internet Routing Registry (IRR) databases. All these sources provide inter-domain (BGP) level informa-

Name	Location	# Peer AS	# Peer Router
Oreg	Oregon, USA	37	45
Eqix	Virginia, USA	4	4
ISC	California, USA	12	14
Linx	London, UK	12	18
Wide	Tokyo, Japan	6	6

Table 2: RouteViews Collectors

tion². Each varies in data content, format, and access method. Some contain both IPv4 and IPv6 information, but we did not use IPv6 information in building the topology.

2.1 BGP Trace Collectors

A BGP trace “collector” is a measurement box that peers with commercial ISP networks via BGP sessions. A collector receives BGP messages from its peers, but it does not advertise any prefixes back to them. Periodically, the collector dumps its full routing tables and routing updates received from its peers.

Table 1 shows a typical entry in a BGP routing table saved by a collector. This entry indicates that AS123, one of the collector’s peers, can reach destination prefix 131.179.0.0/16 via the AS path “123 456 789 987.” From this entry, we determine that the AS-level topology should include four nodes (i.e. 123, 456, 789, and 987), and three links (i.e. 123-456, 456-789, and 789-987).

Typically a collector’s routing table has more than one hundred thousands entries from each peer AS since each peer AS tells the collector how it reaches the entire destination address space. We say that a collector has one “view” of the Internet from each peer AS. The more views (peers) a collector has, the more topological information it can collect.

Besides routing tables, a collector also saves routing updates received from its peers. A routing table shows the preferred paths to reach destination prefixes at a particular moment, while routing updates will reveal alternative paths and backup links over time. Routing updates have a format similar to that of routing tables.

RouteViews and RIPE RIS are two major measurement projects that deploy collectors and make BGP trace data publicly available. Their collectors and peer ASes are listed in Tables 2 and 3. RouteViews has 55 unique views, RIPE has 255 unique views, and combined they have 288 unique views (Table 4). Note that we did not count IPv6 peers in all tables.

²Data from sources such as traceroute collectors can provide detailed information on individual routers and links. These sources are not used since our interest is in establishing the AS-level topology.

Name	Location	# PeerAS	# PeerRouter
RRC00	Amsterdam, Holland	12	12
RRC01	London, UK	51	62
RRC02	Paris, France	27	27
RRC03	Amsterdam, Holland	85	103
RRC04	Geneva, Switzerland	8	8
RRC05	Vienna, Austria	52	62
RRC06	Otemachi, Japan	6	6
RRC07	Stockholm, Sweden	18	18
RRC08	California, USA	4	4
RRC09	Zurich, Switzerland	26	26
RRC10	Milan, Italy	16	17
RRC11	New York, USA	15	15
RRC12	Frankfurt, Germany	23	24

Table 3: RIPE RIS Collectors

	# Peer AS	# Peer Router
RouteViews	55	85
RIPE RIS	255	381
Combined	288	452

Table 4: Unique Peers of RouteViews and RIPE

Overall, RouteViews and RIPE provide both routing tables and updates from their peer ASes and archive past data. We download the data via http, convert the raw binary data to text via *bgpdump* [1], and glean topological information from both routing tables and updates.

2.2 Route Servers

Route servers are routers made publicly accessible by some ISP networks to help troubleshoot network problems. Users can telnet into a route server and run certain router commands. When collecting AS-level topological information from a route server, we run the command “show ip bgp,” which displays the router’s full routing table. Unlike BGP trace collectors, route servers do not provide routing updates, nor do they provide an archive of past data.

From *bgp4.net*, we found 25 route servers in 22 ASes. These route servers peer with more than 43 peers³, of which more than 16 peer ASes are not peers of RouteViews or RIPE (Table 5). In other words, these route servers provide at least 16 more views in addition to the combined routing tables of RouteViews and RIPE.

2.3 Looking Glasses

Looking glasses provide a web interface for running a very limited set of commands on routers. They allow users to check the route to a particular prefix, but do not allow downloading entire routing tables, nor do they provide routing updates. However, some looking glasses allow the command “show ip bgp summary,” whose typical output (from a Cisco router) is like the following:

³We use the command “show ip bgp summary” to determine a route server’s neighboring ASes. Since some route servers do not allow running this command, the total number of peer ASes is underestimated.

Number	#PeerAS	# PeerAS not in RV or RIPE
25	≥ 43	≥ 16

Table 5: Route Servers

Number	# AS	# Router	# AS not in RV or RIPE
174	174	774	132

Table 6: Looking Glasses

BGP router identifier 10.0.0.1, local AS number 123

```
...
Neighbor      V      AS      others
192.168.0.1   4      123     ...
192.168.1.1   4      456     ...
```

We see that the router being queried is in AS123, and it has two neighboring routers, one in the same AS, the other in AS456. Therefore, the AS-level topology should include two nodes (i.e. 123, 456) and one link (123-456). Although a looking glass does not provide its routing table, we can still learn about its direct neighbors. A single looking glass from an ISP usually can query more than one router within the ISP. From *bgp4.net*, we found 174 looking glasses that allow the command “show ip bgp summary,” and they can query 774 routers in total (Table 6). These looking glasses reside in 174 different ASes, and 132 of these ASes do not peer with either RouteViews or RIPE. One looking glass provides much less topological information than one route server, but we have a greater number of looking glasses than route servers, and most of them are not RouteViews or RIPE peers. In the next section we will compare their contributions to the topology.

2.4 Internet Routing Registries

The purpose of the IRR is for operators to coordinate global policy settings. Network operators may register routing policies with the IRR. The databases that form the IRR are manually maintained by operators, mostly on a voluntary basis. Information therein may be incorrect, incomplete, or out-dated. The RIPE portion of the IRR is actively used by ISPs in Europe to filter route announcements and many European exchange points require operators to register with RIPE. Consequently, it is considered the most reliable information in the IRR. We use only the RIPE portion of the IRR for the topology in this paper.

IRR information is expressed in the Routing Policy Specification Language [3], which has 12 different classes of records. We are primarily interested in *aut-num* records, which specify an AS’s import and export routing policies with its neighbors.

```
aut-num:      AS123
as-nam:      ABC
import:      from AS456 accept AS-XYZCUSTOMERS
import:      from AS789 accept ANY
export:      from AS456 announce AS-ABCCUSTOMERS
export:      from AS789 announce ANY
mnt-by:      AS123-MNT
```

Node	RouteViews	RIPE	Route Server	Looking Glass
Individual	18503	18446	18481	2575
Unique	54	9	50	97
Combined	18666			
Link	RouteViews	RIPE	Route Servers	Looking Glasses
Individual	41909	40530	35240	7181
Unique	4071	2684	404	3446
Combined	50010			

Table 7: Topology Snapshots

changed: *contact@xyz.com 20041020*
source: *RIPE*

AS-ABCCUSTOMERS and AS-XYCUSTOMERS are defined in other *as-set* class records. From the import and export policies, we can determine that AS456 and AS789 are neighbors of AS123. The nature of the IRR database is qualitatively different from other data sources, and we are more cautious about adding nodes and links found in the IRR to our topology.

3. TOPOLOGY COLLECTION AND COMPARISON

In this section, we pick a recent date, October 24, 2004, and compile an AS-level topology for this day by utilizing topological information collected from snapshots of routing tables, routing update messages, and data from routing registries. The resulting topology (Section 4) is more complete than the conventional one made by taking routing table snapshot only.

3.1 Snapshots

A snapshot is a way to infer topology from routing tables or looking glass’ neighbor information at a particular time. Barring router mis-configurations and malicious attacks, nodes and links appearing in a snapshot are generally alive at the moment the snapshot is taken.

We take four individual snapshots from four different data sources: routing tables from RouteViews, routing tables from RIPE, routing tables from route servers, and neighbor information from looking glasses. We also make a combined snapshot including all the nodes and links from the four individual ones. On October 24, 2004, we took one routing table from RouteViews, RIPE, and route servers respectively, and queried the looking glasses once, all around the same time. The results are shown in Table 7. The first part of the table describes number of nodes, and the second part of the table describes the number of links. The row of “individual” is the number of nodes or links in a snapshot, the row of “unique” is the number of nodes or links not appearing in the other three individual snapshots. From this table, we make the following observations.

First, no individual snapshot has a significant number of unique nodes. In other words, most ASes show up in all the routing tables. In general, a routing table reflects the reachability to the entire routable IP address space. The reasons that an AS may not show up in a routing table include prefix

aggregation, prefix filtering, multiple origin ASes, and that the AS doesn’t announce a prefix. These cases do happen, but are not common in the current Internet. Therefore, the combined snapshot has only 0.9% more nodes than snapshot obtained from RouteViews routing tables.

Second, there are significant numbers of unique links contributed by individual snapshots. RouteViews peers with 55 ASes, RIPE peers with 255 ASes, but they still each miss thousands AS-level links observed by other snapshots. Although a router peering with n different ASes can learn n different paths to each destination prefix, the router does not learn the information about all the inter-connections among the nodes on these paths. Routing table snapshots taken from different topological views can expose different node connectivities.

Third, RouteViews snapshot exposes more nodes and more links than RIPE snapshot, even though the number of RouteViews’ peer ASes is only about one fifth of RIPE’s. A possible explanation is that most RIPE peers are European ISPs, whose views of the Internet overlap, while RouteViews’ peers spread out over a wider range of topological locations.

Fourth, the route server snapshot shows 35240 links, but only 404 of them are unique, whereas looking glass snapshot has only 7181 links, but 3446 of them are unique. Remember that route servers provide full routing tables, but are small in number (25), while looking glasses provide only neighbor information, but are large in number (174). This suggests that, in terms of collecting topological information, adding more routing tables to existing RouteViews and RIPE data does not help much, because the unique contribution by a router mainly comes from information about the router’s local connectivity. As pointed out by [4], due to the Internet hierarchy and routing policies, most AS paths go up to the top tier ISPs and then go down to the destination ASes. These AS paths overlap after reaching the top tier ISPs. Therefore, the unique contribution by an AS path is mostly about network connectivity before it reaches the top tier ISP.

Overall, compared with RouteViews snapshot, the combined snapshot contains significant more topology information, with 19% more links and 0.9% more nodes.

3.2 Updates

The snapshot of a router’s routing table only captures the topological connectivity contained in the best paths to reach all the destinations at the time of the snapshot. It misses

Node	Combined Snapshot	+Updates(1W)	+Updates(1M)	+Updates(2M)	+Updates(3M)
# Nodes	18666	18756	18917	19072	19261
# Stub Nodes	14402	14466	14601	14738	14901
# Transit Nodes	4196	4222	4248	4266	4291
Link	Combined Snapshot	+Updates(1W)	+Updates(1M)	+Updates(2M)	+Updates(3M)
# Links	50010	52309	55388	57809	60010
# Stub Links	28860	29749	31326	32557	33757
# Transit Links	20814	22224	23725	24914	25904

Table 8: Snapshot plus Updates

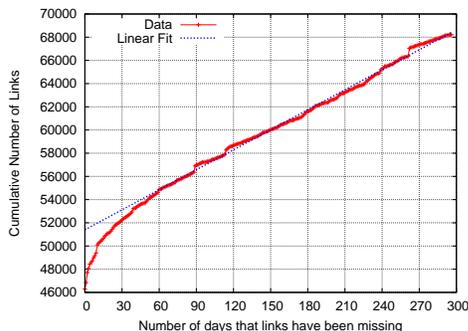


Figure 1: Disappearance Period (links)

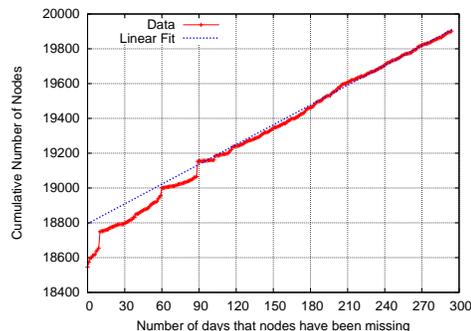


Figure 2: Disappearance Period (nodes)

backup links in the network. The best path may change over time due to link failures, policy changes, or topological changes. Intuitively, collecting BGP routing updates *over time* should reveal backup links that are not included in a snapshot. Using both routing tables and updates collected over a certain time period should reveal nodes and links that occur in any best path seen during the observation period, thus providing a more complete picture of AS connectivity.

Table 8 shows some results of adding topological information from updates to the combined snapshot. The column of “+Updates(1W)” is obtained by adding nodes or links appeared in one week’s updates prior to October 24, 2004. Similarly, the other three columns use routing updates collected during one month, two months, and three months prior to October 24, 2004. We can see that updates provide significant amount of new information.

How long should one process routing updates in order to most efficiently estimate the AS-level topology for a given day (e.g., October 24, 2004)? Table 8 shows that the longer we accumulate the topology from updates, the more links we learn. But while a longer time frame provides more information, it may introduce stale information into the topology as well. Some nodes and links may be removed over time due to business or network operational reasons and by counting these old updates, we may include dead nodes or links. That is, some nodes and links existed at some point of time in the past, but no longer exist in the target day (e.g., October 24, 2004). We should try to maximize the number of live nodes and links collected from the updates while minimizing the number of dead nodes and links.

	Interception a	Slope b	R^2
Link	51396	57.48	0.9952
Node	18795	3.7889	0.9951

Table 9: Results of Linear Regression, $y = a + b * x$

To determine an appropriate time frame to process updates, we study topology accumulation over a longer period, from January 1, 2004 to October 24, 2004, totaling 297 days, using routing tables and updates from both RouteViews and RIPE. We say a node or link has “disappeared” in a day if that node or link does not show up in the routing table or any update within that day. If our study ends on the n^{th} day, a node or link that has not appeared in any routing table or update since day m has a *disappearance period* of $(n - m)$ days. It is the number of days between October 24, 2004 and the day that the node or link is last seen in any routing table and update.

Figures 1 and 2 show the cumulative number of links and nodes over the disappearance period. For example, in Figure 1, a data point at (150, 60000) means that there are 60000 links whose disappearance period is less than or equal to 150 days. In other words, during the 150 days prior to October 24, 2004, 60000 links are observed from RouteViews and RIPE. Both figures exhibit the same pattern: the curve goes up quickly at the beginning, but slows down later on, and stabilizes around a constant increase rate after 2 or 3 months. We use data points between 100 and 300 days to do linear regression, and the results (Table 9) fit the latter part of the curves very well.

The fast increase in the early part of the curve suggests that there are many links and nodes that disappear for a short while, but they are still alive and routing dynamics will likely expose them in updates within one or two months. These are the links and nodes we should include in our topology.

The linear increase in the latter part of the curve suggests that nodes and links that have disappeared for several months are unlikely to come back. If this were not the case, the increase rate should slow down over time, and the curve should eventually flatten out instead of going up at the same rate. Thus it is likely that the nodes and links that have disappeared for a long time no longer exist in the Internet, and we should not include them in the topology. Another dataset, which uses RouteViews routing tables and updates of 17 months (from January 1, 2003 to May 31, 2004), also exhibits the exact same pattern, suggesting that our conclusion is not due to the particular dataset we used here. As part of our future work, we will use other information, such as AS number allocation, routing registry, and WHOIS database, to check whether nodes and links that have disappeared for several months indeed stopped operation.

For the purpose of collecting AS-level topology, Figures 1 and 2 show that using nodes and links appearing within the last 60 days is a good choice. A longer time frame will bring more nodes and links, but the gain will be marginal since most of them may be stale. However, in the topology we publish on the web, all the nodes and links observed since January 1, 2004 will be included, along with timestamps of when it is first observed and when it is last observed. Users can use nodes and links appearing in the last 60 days for general purpose, but they also have the flexibility to adjust the time frame according to their application needs. For example, the timestamp can help the study of topology evolution.

Back in Table 8, we also compare what type of nodes and links different topologies have. We classify nodes and links into two groups: *Stub* and *Transit*. A stub node is one that only appears at the end of an AS path, which means it does not provide transit service to anybody. Other nodes are transit nodes since they appear in the middle of AS paths, thus they provide transit service to someone else. A stub node may have more than one link if it is multi-homed. Most transit nodes are observed both in the middle of some paths and at the end of some other paths because they also originate prefixes. There are only one hundred or so nodes that appear as transit-only. A link connecting a stub node to a transit node is a stub link; a link connecting two transit nodes is a transit link; there is no link connecting two stub nodes. Table 8 shows ⁴ that updates reveal more stub nodes than transit nodes, but reveal about the same amount of new stub links and transit links.

Overall, by adding topology information from two months of routing updates, we improve the combined snapshot by 16% of links and 2% of nodes. Compared with the snapshot from RouteViews routing table only, this is 38% increase in number of links, and 3% increase in number of nodes.

⁴For a small number of nodes and links, we cannot decide their type, since we don't have AS path information about them.

	# records (AS)	# links
all records with links	8097	70222
-void	6177	63625
-void neighbors	4815	23771
-incomplete	3109	13279
-neighbor conflict	2606	7500

Table 10: Validity Analysis of RIPE IRR

3.3 Routing Registries

Due to the manual and voluntary process by which information enters the IRR, we are more cautious about incorporating nodes and links found in the IRR. Our topology only includes links discovered from the RIPE portion of the IRR, which is the most complete and fresh. 8097 (90.4%) out of 8958 ASes that we found in RIPE database registered routing policies that allowed the inference of their links with neighbors. We exclude certain records from being incorporated into our topology if they fail consistency tests. We use heuristics similar to that in [6] to screen the records.

- *Void*: We considered the record for an AS to be void if the AS had not appeared in routing updates or routing tables for more than one month.
- *Void neighbors*: If an AS specifies in its routing policy another AS as neighbor, but we have not seen the neighbor AS in routing updates or routing tables for at least one month, we regard both ASes as void.
- *Incomplete*: We take the last modification time of an AS' record, then find all ASes that appeared as its neighbors in our (snapshot+update) topology since that last modification time. If any of these neighbors do not appear in the said record, we consider this record incomplete because it misses those neighbors.
- *Neighbor conflict*: After passing other checks, if AS A's record shows that AS B is its neighbor, but AS B's record does not show AS A as a neighbor, one or both of these records must be erroneous, incomplete, or stale. Leaning towards caution in accepting information, we discard both.

It is possible that we discard records that contain some valid links just because one or two links in the record are outdated. In the case of conflicting information from possible neighbors, we are unable to determine whether either record is safe to trust. Our goal is to reliably add new links without including erroneous ones. We avoid adding links that we do not consider reliable by trading off the possibility of discarding potentially valid links.

We also risk incorporating stale links from routing updates and even links that only appeared due to mis-configurations. In the IRR, stale and incorrect information will remain in a database until it is removed and thus we are more cautious with IRR data because we cannot place a bound on the staleness of any information. We also err on the side of having a less complete topology over including potentially incorrect links because incompleteness is expected. Despite

Node	Snapshot+Update(2M)	RIPE IRR
Individual	19072	2606
Unique	16466	0
Combined	19072	
Link	Snapshot+Update(2M)	RIPE IRR
Individual	57809	7500
Unique	52804	2495
Combined	60304	

Table 11: Contribution of RIPE IRR

#node	# stub node	# transit node	# unknown type
19072	14738 (77.3%)	4266 (22.4%)	68 (0.3%)
#link	# stub link	# transit link	# unknown type
60304	34235 (56.8%)	25431 (42.2%)	638 (1.0%)

Table 12: The Final Topology on October 24, 2004

our best effort, a record may pass our tests and still contain stale information. In the published topologies, we label each node and link whether it is only observed from IRR or not. If the use of the topology is sensitive to even a small amount of stale information, users can discard nodes and links that only come from IRR.

The number of nodes and links after passing each test is listed in Table 10. Nodes and Links that pass all the tests are combined with “Snapshot + Updates(2M)” in Table 11.

4. THE FINAL TOPOLOGY

After incorporating topological information from all routing tables, router neighbor information, routing updates, and registry, we obtained a final topology of 60304 links and 19072 nodes as the AS-level Internet topology on October 24, 2004. From Table 12, we can see that transit nodes are only about one fifth of all the nodes in the Internet, but over 40% of links connect transit nodes to other transit nodes. In the topology, the average degree of transit nodes is 19.9, and the average degree of stub nodes is 2.3. Compared with a topology map built from conventional RouteViews snapshot, our topology has 44% more links and 3% more nodes. Perhaps more importantly, our more complete topology is more densely connected than the one derived from only the RouteViews snapshot.

To provide some initial characterization of the topology, we first examine node degree distribution. Figure 3 shows the node degree distribution for four different intermediate topologies that are constructed as we build toward a final topology. The X -axis is node degree, the Y -axis is the complementary cumulative distribution function of node degree, and both axes are in logarithmic scale. As the intermediate topologies become more complete (moving from a single snapshot to combined snapshot plus two months of updates and IRR data), the percentage of nodes with degree between 3 and 400 increases. But in general, these curves share same basic characteristics and can be well fitted by linear regression, often referred to as the “power-law” of Internet topology. For instance, the linear fit for the final topology has $R^2 = 0.976$, and it gives the power exponent as -1.166 .

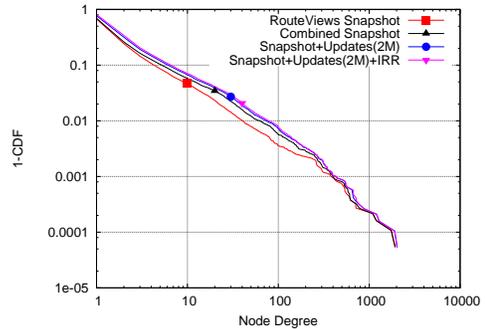


Figure 3: Node Degree Distribution

In addition to simple node degree distribution, we also examined the impact of commercial relationships on the graph structure. The Internet has a hierarchical structure that is determined by the commercial relationships between ASes. At the top of the hierarchy are a small number of “tier-1” providers, who are peers to each other, and provide transit service to the rest of the Internet. At the bottom of the hierarchy are thousands of stub ASes, who do not provide transit service to anyone else. In the middle are non-tier-1 transit ASes, who are customers to upper level ASes, but providers to lower level ASes. To calculate node degree distribution for these three types of nodes, we need first separate tier-1 transit ASes from non-tier-1 transit ASes.

A tier-1 AS does not have any provider and should peer with all other tier-1 ASes in order to reach all destination networks in the Internet. As a result, tier-1 ASes together should form a clique in the topology graph. Tier-1 ASes typically have many customers, including both stub ASes and lower level service providers. We divide a node’s degree into two parts: *transit degree*, which is the number of links to transit nodes and *stub degree*, which is the number of links to stub nodes. Compared with non-tier-1 transit ASes, tier-1 ASes should generally have larger transit degree.

Ideally, we could identify the tier-1 ASes by examining Internet peering relationships (e.g., provider, customer, peer). However, peering relationship data usually is not publicly available and is instead inferred from BGP data. The inference methods can introduce errors. In fact, one of the latest relationship inference algorithms [19] uses a list of known tier-1 ASes as the *input* to improve the inference result. We use a heuristic that is based on the assumption that tier-1 ASes should form a clique and should have large transit degree. Initially, the node with highest transit degree is counted as a tier-1 AS. Then we consider other nodes one by one following descending order of transit degree. A node is added into tier-1 if it has links to all existing tier-1 nodes, or misses at most one link to existing tier-1 nodes. The final result depends on the order of nodes. We follow descending order of transit degree to favor nodes with high transit degree, and allow one missing link to accommodate possible incompleteness of our topology. This heuristic is similar to the one in [10], but [10] uses the total node degree rather than transit degree.

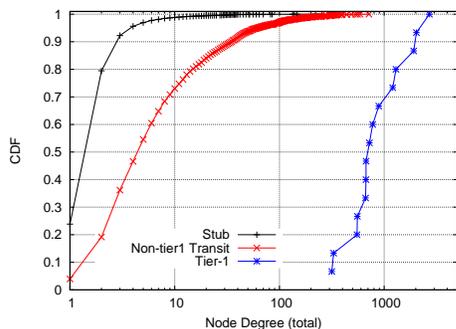


Figure 4: Total Degree Distribution

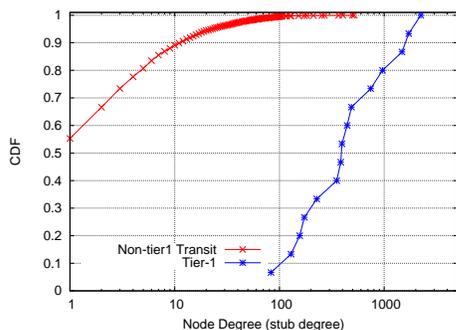


Figure 5: Stub Degree Distribution

Using this method, we identified 15 nodes as the tier-1 nodes. These nodes nearly form a clique, missing only 8 links. Distributions of different degrees and different types of nodes are shown in Figures 4, 5, and 6.

In Figure 4, some non-tier-1 transit ASes have total degree of 1. At first glance, a node with degree one would seem to be a stub AS. This is because our topology only includes data from two months before October 24, 2004. However, the classification into transit and stub nodes is based on a ten-month dataset starting from January 1, 2004. We use two month’s data to get “live” links, but take the ten month’s data to determine whether a node is transit or not. The transit nodes with degree of 1 in the two-month topology actually have more than one link in the ten-month topology. This probably means that some links are missed in the two-month topology rather than the node has changed from transit to stub. Our justification is that links may come and go relatively frequently, but the type of a node depends on whether a real world business entity is an ISP or not, which, if ever to change, should change very infrequently.

These degree distributions show that, in general, lower level nodes have fewer links compared with upper level nodes, but a small number of lower level nodes can have large degree. This is consistent with the hierarchical classification results in [18]. We can also see that the variance in tier-1 nodes’ total degree largely comes from the variance in their stub degree; their transit degree is within a much narrower range.

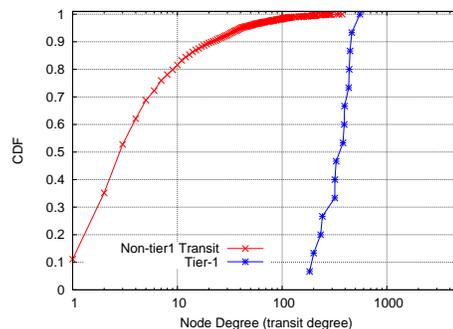


Figure 6: Transit Degree Distribution

The results here have focused on the topology generated on October 24th, 2004. However, it is important to note that we have automated the data collection process and topology generation and continue to generate new topologies each day. Data is downloaded from RouteViews, RIPE, route servers, and looking glasses every night⁵. A new topology containing nodes and links from all sources is generated and published on the web (<http://irl.cs.ucla.edu/topology/>) every day. Each day’s topology has two files, a node file and a link file. Each line in a node file lists the AS number of a node, and each line in the link file lists the two incident AS numbers of a link.

Since not all researchers seek the same topological information, three types of auxiliary information are also included in the files: (1) timestamp, specifying when a node or a link is first seen, and when it is last seen, (2) type, specifying whether a node or a link is stub or transit, and (3) data source, whether a node or a link is only seen from the routing registry data or not. This information allows users to make better use out of the topology. For example, if a user is only interested in the transit connectivity of the Internet, they can easily get it from our topology files by using only transit nodes and links.

5. RELATED WORK

In an early work Govindan *et al.* [11] derived the AS topology from routing updates collected over 21 days. This was one of the first work in this area, however the methodology of accumulating topological information from updates was not studied, probably due to unavailability of RouteViews and RIPE RIS data at that time.

Despite popular usage of AS-level topology in a wide range of research activities, the methodology of inferring such topology has not been carefully studied, nor has an up-to-date topology been maintained and made readily available. Chang *et al.* [6] examined the completeness of connectivity information collected from RouteViews routing table snapshots. They found that Tier-1 ASes are well covered in the snapshot, while ASes at lower tiers have poor coverage. As a result, the snapshot misses a significant amount of connectivity information. They obtained additional AS connec-

⁵Since fetching full routing table puts load on the router, we sent emails to route server administrators asking for permission before starting the daily download.

tivity information from route servers, looking glasses and routing registry databases. Our work explores an additional dimension of accumulating information from updates over time. We believe it is important to both collect from more data sources and accumulate over time in order to obtain more complete AS connectivity information.

Lakhina *et al.* [12] studied the “(k-m)-traceroute” method for its sampling bias in finding node degree distribution. Assuming a topology G of n nodes, k source nodes and m destination nodes, and that shortest paths between every source-destination pair are known (e.g., by traceroute probing or BGP routing table), then the union of all these shortest paths is the observed topology G' . They showed that G' could have a very different node degree distribution from G , because nodes with higher degree are sampled more frequently than nodes with lower degree. Barford *et al.* [4] studied the marginal gain of adding additional source or destination nodes in the same model. They found that a small number of source nodes are necessary, but adding more does not gain much. On the other hand, adding more destination nodes gives a linear increase in useful information. These conclusions in [12] and [4] do not directly apply to our work because (1) in BGP almost all nodes appear as destinations, i.e., $m \approx n$, which is quite different from traceroute data, and (2) we collect connectivity information from routing updates over time, instead of just a static snapshot. An interesting future work would involve applying their analysis to our data to check potential sampling bias and to study the marginal gain of adding additional monitoring peers.

In [2], CAIDA uses another method to derive the AS-level topology. They conduct traceroute from 25 sources to hundreds of thousands of destinations to gather IP addresses of intermediate routers. Then they map each router to its residing AS, thus deriving nodes and links in the AS-level topology. As a future work, we will look into this method to see if it can make a unique contribution to our final topology.

6. SUMMARY

The contribution of this work is three-fold. First, we assembled the most complete AS-level topology from as many inter-domain routing sources as we can get, including RouteViews, RIPE, route servers, looking glasses, routing registry, and routing updates. Compared with the routing table snapshots taken from RouteViews, which is commonly used in research, our topology on a sample day contains 44% more links and 3% more nodes. Second, we accumulate topological information from routing updates over time, and develop a method of determining a good disappearance period to use. Third, we make the topology publicly available and will keep it up-to-date. Given the wide usage of AS-level topology, we expect this service to benefit the research community as a whole. As for future work, we will explore new methods of collecting AS-level nodes and links, including using traceroute results and improving the screening of routing registry data. We will also explore other auxiliary information that may facilitate usage of the topology. For example, it may be useful to include in the topology AS peering relationships and routing policy inferred by well-known algorithms.

7. REFERENCES

- [1] bgpdump.
<http://www.ris.ripe.net/source/libbgpdump-1.4.tar.gz>.
- [2] CAIDA’s macroscopic topology AS adjacencies.
http://www.caida.org/tools/measurement/skitter/as_adjacencies.xml, June 2004.
- [3] C. Alaettinoglu, C. Villamizar, E. Gerich, D. Kessens, D. Meyer, T. Bates, D. Karrenberg, and M. Terpstra. Routing policy specification language (RPSL). RFC 2622, Internet Engineering Task Force, June 1999.
- [4] P. Barford, A. Bestavros, J. Byers, and M. Crovella. On the marginal utility of network topology measurements. In *ACM SIGCOMM Internet Measurement Workshop (IMW)*, 2001.
- [5] G. D. Battista, M. Patrignani, and M. Pizzonia. Computing the types of the relationships between autonomous systems. In *Proc. of IEEE INFOCOM*, Mar. 2003.
- [6] H. Chang, R. Govindan, S. Jamin, S. J. Shenker, and W. Willinger. Towards capturing representative AS-level Internet topologies. *Computer Networks*, 44(6):737–755, Apr. 2004.
- [7] Q. Chen, H. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger. The origin of power-laws in Internet topologies revisited. In *Proc. of IEEE INFOCOM*, June 2002.
- [8] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the Internet topology. In *Proc. of ACM SIGCOMM*, 1999.
- [9] L. Gao. On inferring autonomous system relationships in the Internet. *ACM/IEEE Transactions on Networking*, 9(6):733–745, 2001.
- [10] Z. Ge, D. Figueiredo, S. Jaiwal, and L. Gao. On the hierarchical structure of the logical Internet graph. In *ITCOM*, 2001.
- [11] R. Govindan and A. Reddy. An analysis of Internet inter-domain topology and route stability. In *Proc. of IEEE INFOCOM*, 1997.
- [12] A. Lakhina, J. W. Byers, M. Crovella, and P. Xie. Sampling biases in IP topology measurements. In *Proc. of IEEE INFOCOM*, 2003.
- [13] A. Medina, A. Lakhina, I. Matta, and J. Byers. BRITE: An approach to universal topology generation. In *Proc. of MASCOTS*, Aug. 2001.
- [14] K. Park and H. Lee. On the effectiveness of route-based packet filtering for distributed DoS attack prevention in power-law Internets. In *Proc. of ACM SIGCOMM*, pages 15–26, 2001.
- [15] Y. Rekhter and T. Li. Border Gateway Protocol 4. RFC 1771, Internet Engineering Task Force, July 1995.
- [16] The RIPE Routing Information Services.
<http://www.ris.ripe.net>.
- [17] The Route Views Project.
<http://www.antc.uoregon.edu/route-views/>.
- [18] L. Subramanian, S. Agarwal, J. Rexford, and R. H. Katz. Characterizing the Internet hierarchy from multiple vantage points. In *Proc. of IEEE INFOCOM*, June 2002.
- [19] J. Xia and L. Gao. On the evaluation of AS relationship inferences. In *Proc. of IEEE GLOBECOM*, Dec. 2004.