# Statistical implications of augmenting a BGP-inferred AS-level topology with traceroute-based inferences

Bradley Huffaker, Marina Fomenkov, kc claffy
{bradley,marina,kc}@caida.org
CAIDA, University of California, San Diego

## ABSTRACT

The limitations of a BGP-inferred AS-level topology are generally understood, but the impact of various types of missing links on various topology properties and characteristics remains an open question. As part of CAIDA's continuing work to improve the completeness, accuracy, and richness of the measured AS-level Internet graphs, we developed a methodology to combine different types of data into a comprehensive ("**combined**") Internet topology.

Our methodology has three steps: deriving a base graph from AS paths observed in publicly available BGP by breaking AS paths into AS links; augmenting this base graph with traceroute-derived inferred AS links (corresponding to only the first AS hop in the traceroute, for methodological reasons); and extracting AS level topology data from multilateral peering registration information in European Internet eXchange Point (IXP) route server data. Specifically, we examined how the introduction of 241,459 additional peering links (a 136% increase over the BGP graph) and 144 additional nodes (0.3% increase) inferred from traceroute and Internet eXchange (IX) data changed the topological properties of the AS-level graph originally derived from the BGP data. Notably, only 6.8% of the ASes in the original BGP-based graph gained additional links. Of those ASes, links were primarily added to medium degree ASes, and ASes classified as edge ASes remained largely unaffected. For all four metrics we used to define the peripheral part of the graphs (customer cone, coreness, eccentricity, node betweenness), the change of the relative size of this part was 3% or less between the BGP-based and the combined graphs (see Table 6). One of the primary insights of this exercise is that for the largest and smallest degree ASes, BGP measurements capture connectivity characteristics well, but for many middle-degree ASes, additional connectivity not visible in global BGP data repositories is revealed by traceroute-based inferences.

## Keywords

Internet topology, active measurement, passive measurement, AS

## 1. MOTIVATION

Topology maps of the Internet are indispensable for characterizing this critical infrastructure and understanding its properties, dynamics, and evolution. Most Internet mapping methods have focused on characterizing and modeling the network structure at the level of interconnected Autonomous Systems (ASes). AS-level topologies (or graphs) annotated with business relationships between the ASes provide insights into technical, economic, policy, and security needs of the growing and evolving Internet ecosystem. As part of CAIDA's continuing work to improve the completeness, accuracy, and richness of the measured AS-level Internet graphs, we developed a methodology to combine data from three different types of sources – BGP data, traceroute-inferred first hop AS links, and IXP-peering information from registration databases, into a comprehensive ("**combined**") Internet topology. This document presents this methodology for creating a combined topology graph, as we applied it to these three different sets of macroscopic Internet topology data gathered in 2014. We then quantitatively compare the BGP-only graph that we previously used, to this larger graph that leverages other sources of information. We perform this comparison in terms of basic graph-theoretic metrics, reflecting both the characteristics of links and nodes in each graph as well as each graph as a whole.

| graphs | links | | | nodes |
|---|---|---|---|---|
| | c/p | peer | total | total |
| **BGP-derived** | 93,539 | 83,852 | 177,391 | 46,177 |
| traceroute-derived | 0 | 7,166 | 7,166 | 2,432 |
| IX-derived | 0 | 264,803 | 264,803 | 1,555 |
| **combined** | 93,539 | 325,312 | 418,851 | 46,321 |

**Table 1:** ***Number of nodes and links derived from various data sources. The numbers of nodes and links in the combined graph are fewer then the sum of those values in the partial graphs due to overlap across data sources.***

This exercise provides qualitative as well as quantitative insights into the value of the additional data sources relative to the canonical approach to AS-level topology inference and analysis, and we offer guidelines for how to apply these insights in practice. For example, we found that the peering inferences from IXP data tend to enrich connectivity between nodes with moderate degrees (i.e., neither well-connected hubs, nor nodes with few neighbors). This result is consistent with the means of acquiring the data, and of the economics of the Internet peering ecosystem. That is, the data comes directly from IXP route server databases, which tend to reflect a rich mesh of multilateral peering among nodes present at that IXP, rather than transit relationships between a small customer and a larger provider. The larger providers tend to not peer directly with moderately sized (or smaller) networks; instead they seek to cultivate such networks as customers. Similarly, the smaller networks tend to procure transit services from large ISPs until they are large enough themselves to negotiate peering relationships with other moderately sized (or larger) networks.

## 2. METHODOLOGY

Our methodology involves the following three steps. First, we derive a base graph from AS paths observed in publicly available BGP data [1, 2] by breaking AS paths into AS links [3]. A ma-

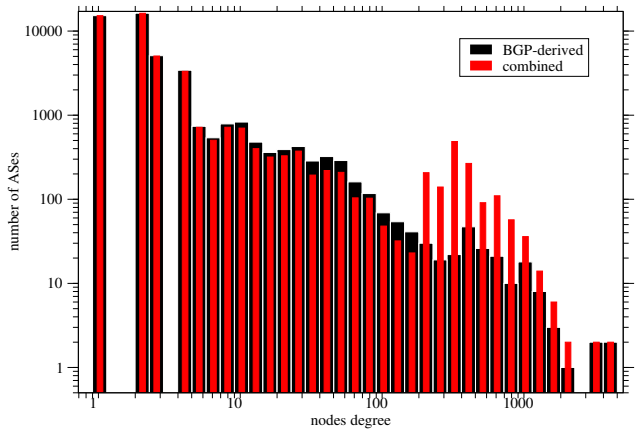


(a) The AS degree distribution roughly follows a power law in the BGP-derived graph. The addition of 241,460 new peering links derived from the traceroutes and IX data noticeably changes this behavior: the number of ASes with degrees between 20 and 200 decreased, while the number of ASes with degrees $> 200$ increased.

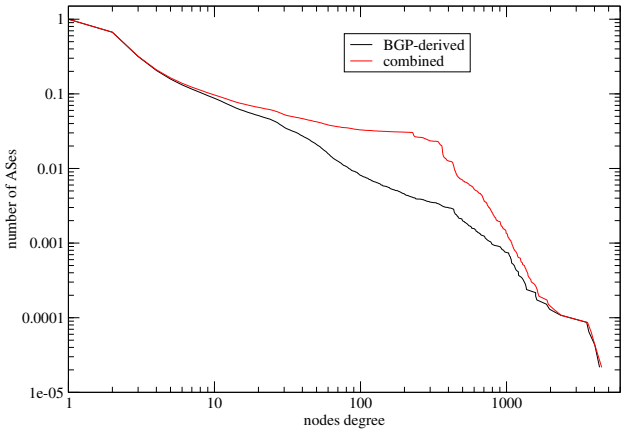


(b) The CCDF of the AS degree distribution. The additional peering links increase the percentage of ASes with degrees more then 230 from 0.4% to 2.9% of nodes. This change is an observable manifestation of the "flattening" of the Internet AS topology due to an increasing number of peering connections.

**Figure 1: AS degree distributions for December 2014 data set (Section 3.1.1).**

jor limitation of such a BGP-derived graph stems from the fact that any given BGP monitor sees only a small fraction of the overall topology. One reason for this partial view is that BGP specifies that at each hop, routers must forward only the optimal paths based on local preferences, rather than announcing all paths. Business relationships between ASes also limit the set of forwarded paths: while an AS's customers will potentially see all of its links, its peers and providers will only see its customer links. Thus, publicly available BGP data (Route Views and RIPE RIS) provide good coverage of customer-provider links, but are less suitable for observing peering links: a BGP-derived graph includes peering links immediately connected to either a provider of a BGP collection point, or this provider's provider, but in general it substantially undersamples peering links. To improve peering link coverage, we have examined two additional data sources: traceroute data continuously collected by our Ark measurement infrastructure [4] and routing relationship data registered with public route servers at Internet eXchange Points (IX) [5].

The second step of our graph-constructing methodology is to process traceroute data collected by Ark monitors to discover previously unobserved links (i.e., those not visible in BGP) that involve each Ark monitor's hosting AS and this AS's neighbors. Since customer-provider links are generally visible in BGP, we infer these links to be peering (i.e., not transit).

A classic problem with deriving AS links from traceroute measurements is that the obvious approach – mapping IP addresses to their origin ASes and thus converting IP links to AS links – is prone to inferring false peering links with neighbors of actual peers. One cause of such misinference happens when multiple ASes attach to the same router and the router responds to a TTL-expired message with an outbound interface address on the forwarding path rather than the inbound interface on the receive path. Another misinference risk is the *missing link* problem, where the only address observed on customer Y's router comes from provider X's router, and subsequent in the traceroute path we see an interface from Z, a neighbor of Y, therefore inferring a link between X and Z where none exists.

To mitigate this problem, we use Luckie, *et al.*'s methodology [6] to collapse observed interfaces into routers and infer which AS owns each router. The ownership findings are based on IP-to-AS mapping derived from public BGP data (using [7]), list of peering prefixes from *PeeringDB* [8], and CAIDA's inferred business AS relationships [3]. Then we convert the observed IP path into an AS path using the router ownership information (rather than mapping each observed IP to AS directly) and add the first interdomain AS link in the resulting path to the AS graph. Accurately inferring subsequent interdomain links observed in traceroute requires overcoming the natural sample bias in traceroute [9], i.e., poorer visibility into distant networks, which limits our ability to assemble constraints. To validate this approach, we compared the inferred peering links with observed ones for 17 Ark monitors that were co-located with BGP monitors in October 2014. We found that 97% of the traceroute-inferred peering links were visible in BGP. We expect a similar level of coverage for peering link inferences we made for Ark monitors not co-located with BGP monitors.

In our third step, we examine data from 13 large European Internet eXchange Points (IXPs). By juxtaposing route server data with a mapping of BGP community values, we infer peering links that cross those exchanges: the IX-derived or *multilateral* peering [10].

Here we report on our analysis and curation of BGP, traceroute, and IX datasets collected in December 2014, in which we merged all obtained links to create a new *combined* topology. Table 1 shows the basic statistics of each partial and the total combined AS graphs. Compared to the basic BGP-derived graph, the combined graph grew from 46,177 ASes with 177,391 links to 46,321 ASes with 418,850 links. We assume that the BGP data captures all customer-provider links [11], which implies that the two graphs have the same number of customer/provider links (in this case 93,539), while the additional data we used to augment the graph contributes only new peering links – in this case the number of peering links, nearly tripled increasing from 83,868 to 241,459. Note that we have no access to traffic data so we are not in a position to judge how important these links are to global traffic flow.

## 3. TOPOLOGY CHARACTERISTICS

We compared our new combined AS-level topology with a straight BGP-based topology using a set of metrics described in [12]. Our tool *topostats* [13] calculates these metrics under a strictly geo-

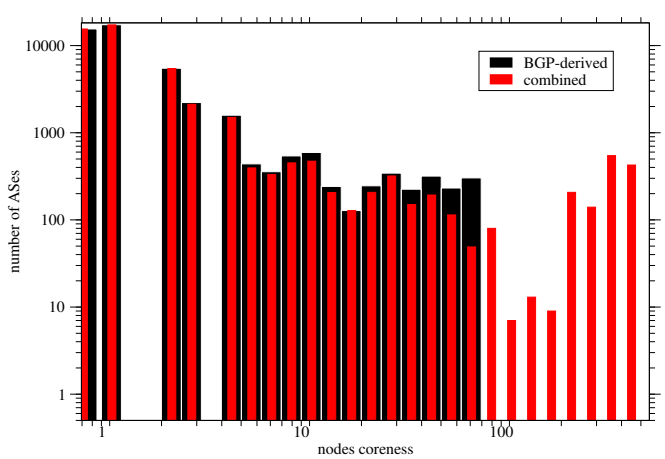


(a) The coreness distribution. The leftmost column includes ASes with coreness of 0. For small values of coreness ($< 22$), the distribution follows a power law and is similar in both graphs. Due to additional links in the combined graph, the coreness of 3.2% of ASes exceeded the previous maximum of 76 in the BGP-derived graph, reaching the new maximum of 409.

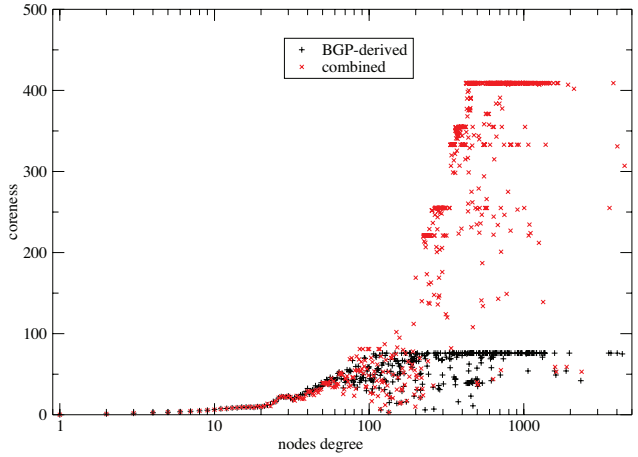


(b) The average coreness as function of node degree. Coreness values are largely unchanged for ASes with degrees $< 100$. For larger ASes, the coreness increased, raising the maximum value observed in the graph from 76 to 409.

**Figure 2: Coreness, i.e., the maximum k such that the node remains the k-core but is removed from the k + 1-core, for December 2014 data set (Section 3.1.3).**

metric interpretation of the graphs: undirected, with policy-free routing via the shortest paths. We subdivide our metrics into two categories: metrics describing local connectivity in a graph (Section 3.1) and characteristics of the global structure of the topology (Section 3.2). Table 2 summarizes all calculated metrics for the BGP-based and the combined graphs.

## 3.1 Local connectivity

This section describes the results of computation of three types of metrics that describe local characteristics of the graph – degree, clustering, and coreness – and their implications.

### *3.1.1 Degree*

An AS's ***degree*** is *the number of nodes with a direct link to this AS in the AS graph.* The average AS degree is $\bar{k} = 2m/n$, where $n$ is the number of ASes (nodes) and $m$ is the number of links in the graph. Adding the 241,460 new links derived from traceroute and IX data more than doubled the average AS degree, increasing it from 7.7 to 18.1. The maximum observed degree increased from 4,333 to 4,484.

Figure 1(a) presents the ***AS degree distributions*** for both (the BGP-based and combined) graphs using a log-log scale with degrees (the $x$-axis) binned into 10 bins per decade. Figure 1(b) plots the CCDF of the AS degree distribution on a log-log scale. Comparison of the degree distributions for the two graphs shows that the additional links in the combined graph do not follow a power law distribution for ASes with degrees greater than 10. The number of ASes with degrees less than 10 did not change, the number of ASes with degrees between 10 and 230 decreased; and the number of ASes with degrees greater than 230 increased from 0.5% of the total for the BGP-derived graph to 2.9% of the total for the combined graph. These changes in the long tail of the degree distribution represent the so-called "flattening" of the Internet AS topology: an increasing number of peering links that has altered the previous model of hierarchical AS connectivity best approximated by a power law of the degree distribution.

Another degree-related metric is ***the average neighbor connectivity (ANC)*** $k_{nn}(k)$ defined as *the average degree of the neighbors of nodes of degree* $k$. Table 2 shows the average values of this metric, $\bar{k_{nn}}$, that is, ANC averaged over all degree values in the graph. For comparison between graphs of different sizes, we normalize ANC by its maximal possible value $n - 1$ (reachable in full mesh graphs). The normalized average ANC is only 9% higher in the combined graph than in the BGP-derived one; the difference is small because there is no noticeable increase in degrees for most ASes (those with degree $< 10$, which comprise 91% and 90% of the total for the BGP and the combined graph, respectively).

### *3.1.2 Clustering*

***Local clustering*** is defined as *the fraction of an AS's neighbors connected by a direct link.* A clustering value of 0 means that none of an AS's neighbors connect to each other, while a value of 1 means that they all interconnect.

Figure 3(a) shows the local clustering distribution for our graphs, plotted using linear bins (0.02 wide) on the $x$-axis and log scale on the $y$-axis. The distribution is almost uniform across the $[0, 1]$ range of possible local clustering values except for the large peaks at the ends of the interval. The addition of peering links in the combined graph increased the fraction of ASes with a local clustering value of 1 from 18% to 22% of the total, while the fraction of ASes with a local clustering value of 0 decreased from 59% to 55%. Figure 3(b) shows that for both graphs, local clustering is almost constant for ASes with degrees $< 100$ and generally follows a power law for larger degrees, with exponent -0.86. High values of local clustering for medium-degree ASes ($200 < k < 500$) in the combined graph reflect the fact that these ASes tend to peer with each other. Peering links tend to increase the robustness of connectivity for those nodes: the more interconnected are the neighbors of a given node, the higher path diversity around the node. (Section 3.2.2 describes how the increase in available path diversity affects another topological metric: link betweenness.)

The mean local clustering metric shown in Table 2 rose from 0.28 to 0.34. Table 2 also includes the ***clustering coefficient*** defined as *the percentage of 3-cycles among all connected AS triplets in the entire graph*; it increased almost 10-fold from 0.054 to 0.47. The drastic difference in the impact of additional links on the mean

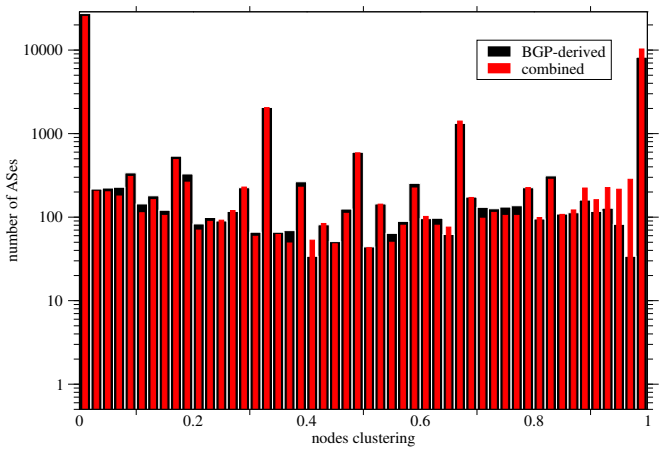


(a) Local clustering distribution. Except for the peaks at the extreme values of 0 and 1, ASes appear to be uniformly distributed across possible values of local clustering.

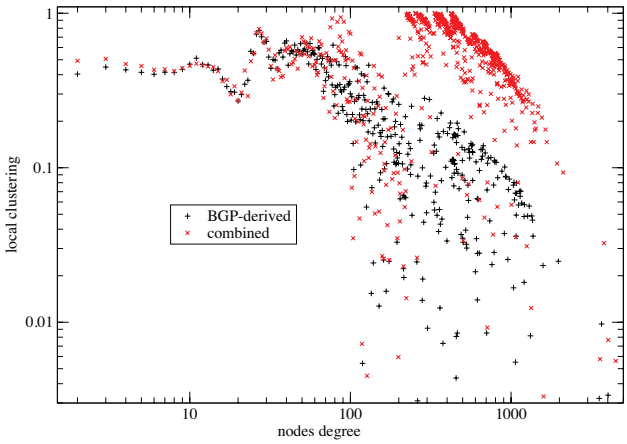


(b) Average local clustering as function of AS degree. For small degrees $< 70$, the average local clustering is nearly constant and remains mostly unchanged in the combined graph. For larger degrees the average local clustering is decreasing as power law with exponent -0.86. In the combined graph, the local clustering noticeably increased for ASes with degrees $> 200$.

**Figure 3: Local clustering (fraction of an AS's neighbors that are directly connected) for December 2014 data set (Section 3.1.2).**

local clustering and clustering coefficient stems from the difference in how these two metrics weight nodes. The mean averages across all nodes, but nodes with degrees $< 60$ comprise about 98% of both graphs; Figure 3(b) shows that for those lower-degree nodes, clustering remained essentially the same – thus, there is only a 21% increase in mean local clustering in the combined graph. In contrast, the clustering coefficient counts node triplets, and since large degree nodes participate in many more triplets than small degree nodes, they contribute more to this metric. Figure 3(b) shows a large increase in clustering for nodes with degrees $> 200$, which explains the large increase in the clustering coefficient.

### *3.1.3 Coreness*

An AS graph's $k$-core is the subgraph obtained from the original graph by iteratively removing all nodes of degree $k$ or less [12]: An AS's ***coreness*** $\kappa$ is *the maximum $k$ such that the node remains in the $k$-core but is removed from the $k+1$-core.*

Figure 2(a) plots the distribution of coreness in both graphs using a log-log scale with coreness values (the $x$-axis) binned into 10 bins per decade. Edge ASes of degree 1, as well as ASes of degree 2 connecting them to the main graph, all have coreness 0 (the leftmost column). In both graphs there are slightly more nodes with coreness of 1 than 0. Additional links in the combined topology caused a significant increase of its coreness: the average coreness increased from 3.0 to 11.9, and the maximum coreness of the graph increased from 76 to 409, with 3.2% of the ASes in the combined graph having coreness values $> 76$.

Figure 2(b) shows the average coreness as function of node degree. Horizontal groups of dots indicate peaks in the coreness distribution at values of 221, 255, 333, 355, and 409. These clusters do not necessarily mean that many ASes with these degree values exist. More likely, these peaks reflect a single AS with that degree acting as a lynchpin for a cluster of interconnected ASes. Thus, all ASes in this cluster have the same coreness since removal of the lynchpin AS from the graph causes the removal of the whole cluster.

## 3.2 Global structure

### *3.2.1 Distance*

The first global metric we examine is the ***shortest path length*** or the ***distance*** between ASes. For each AS, we find the average distance (AD) between this AS and all other ASes in the graph. Figure 4(a) (plotted in linear coordinates on both axes) shows that the average distance distribution changed from a bipolar one with maxima at 3.1 and 3.6 in the BGP-based graph to a tripolar one in the combined graph, with maxima at 3.1, 3.5, and 4.0. Although the combined graph has 2.4 times more links than the BGP-based one, the overall average path length for the whole graph decreased by only 1.9% (from 3.75 to 3.68). This insignificant decrease is unsurprising, since most links added to the combined graph rely on inferences from IXP data, which reflect pairs of ASes that are already near each other, so that the new links do not create shortcuts across the longest paths. Additionally, most paths in the graphs are already rather short, and there is not much room for further shortening.

Next, we calculate the **average average distance (AAD)** for ASes of a given degree. Figure 4(b) (plotted in log-linear coordinates) shows that for both graphs, AAD behaves as a power law function of the node degree with exponent of $-0.1$ in practically the whole range of observed degrees. Figure 4(b) also illustrates how additional links in the combined graph affect the AAD. For ASes with small degrees ($< 40$), their average distance to other ASes in the graph decreased, because for those ASes even a few extra links, while not increasing the degree by much, may drastically influence their connectivity (for example, if an added link connects a small degree AS to a hub of ASes well connected to others). At the same time, the AAD increased for medium degree ASes ($40 < k < 250$) while for large degree ASes the AAD distributions overlap. Medium and large degree ASes are already well connected; therefore, their degree grows due to the addition of new links, but the average distance remains the same.

Another distance-related graph characteristic is ***AS eccentricity*** defined as *the length of the longest shortest path for this AS*. The

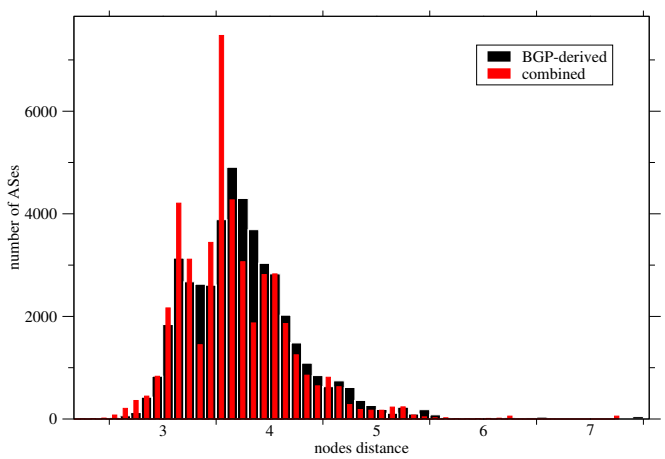


(a) The average distance distribution. For each AS, we calculate the average distance from this AS to all other ASes in the graph. The distribution changed from bipolar to tripolar while the average path length decreased from 3.75 down to 3.68.

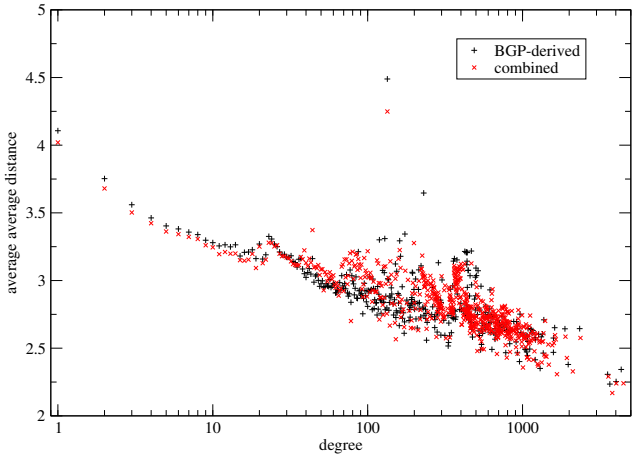


(b) We calculated the average average distance for all ASes of the same degree. This distance metric is a power law function of node degree with the exponent of $-0.1$ for both graphs.

**Figure 4: Average distance between a given AS and all other ASes in the graph for December 2014 data set (Section 3.2.1).**

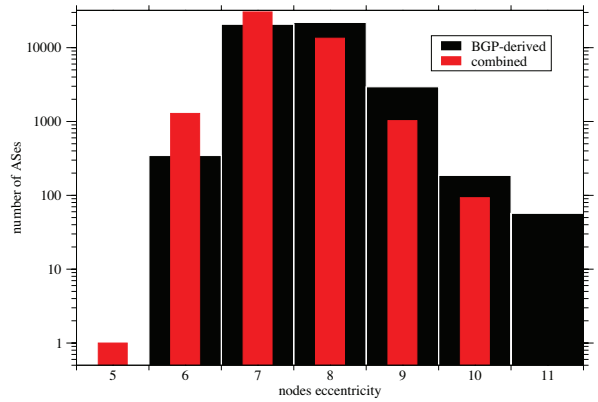


**Figure 5: AS node eccentricity distribution for December 2014 data set. The additional links shifted the distribution toward smaller values, with average eccentricity decreasing from 7.6 in the BGP graph to 7.3 in the combined one. The eccentricity of AS 1299 (Telia) changed from 6 to 5. Telia is the largest European AS, and the new links in the combined graph derived primarily from European IXP-based peering adjacencies.**

largest eccentricity value is the graph's *diameter* and the smallest eccentricity value is the graph's *radius*. Both of these characteristics decremented by one, from 11 and 6 in the BGP graph, to 10 and 5, respectively, in the combined one. The average eccentricity (averaged over all ASes in the graph) decreased by 3.9%, from 7.6 to 7.3.

### 3.2.2 *Betweenness*

***Betweenness*** is *the number of shortest paths that cross a given AS (node) or a link in the graph normalized by its maximum possible value of* $n(n-1)$ [12]. This metric reflects how central a given node or link is to the graph. Elements near a bottleneck or the center of a cluster have higher betweenness values. Elements peripheral to the topology have low values. Most ASes have betweenness 0 (consistent with most ASes located at the edge). We calculate node and link betweenness using a fast algorithm from [14]. Figure 6 shows the node and link betweenness distributions in log-log scale with betweenness values (the $x$-axis) binned into 10 bins per decade.

The node betweenness distributions (Figure 6(a)) for both graphs are similar. The maximum node betweenness increased by 6.7% (from 0.134 to 0.126), but the average node betweenness decreased by 3%, from $6.0 \times 10^{-5}$ down to $5.8 \times 10^{-5}$.

Figure 6(b) shows that in contrast to similar node betweenness distributions, the link betweenness distributions are notably different. There is a substantial increase in the number of links with small values of link betweenness in the combined graph. As the result, the average link betweenness dropped by 58%. The maximum link betweenness also decreased, from 0.0084 to 0.0071 (by 15%). The metric of betweenness in some sense reflects the *centrality* or *importance* of an element to the graph. Clearly, the addition of the new links in the combined graph decreased the relative importance of any given link while the importance of each AS did not change.

Figure 7 shows that in the BGP-based graph, the node betweenness is a growing power-law function of node degree with the exponent value of 1.034. Addition of a large number of peering links in the combined graph breaks the power law of the degree distribution for some node degrees, i.e., the degree of some nodes became much larger, but they continue to conduct the same fraction of the shortest paths, and therefore their betweenness did not change.

## 4. ROUTING POLICIES AND TOPOLOGY

In calculating the topological characteristics discussed in Section 3, we invoked strictly geometric interpretation of the graphs assuming policy-free routing via the shortest paths. Yet routing in the Internet is sometimes constrained by complex interdomain policies. In this section we compare various characteristics of the two graphs that take into account the business relationships between ASes inferred according to CAIDA's AS Ranking algorithm [3], and the impact of these relationships on graph properties.

### 4.1 Customer and Peer Cones

These two metrics, customer and peer cones, attempt to quantify policy-compliant topological characteristics. For a given AS $A$, the ***customer cone*** is the *set of ASes that pay it, directly or indirectly,*

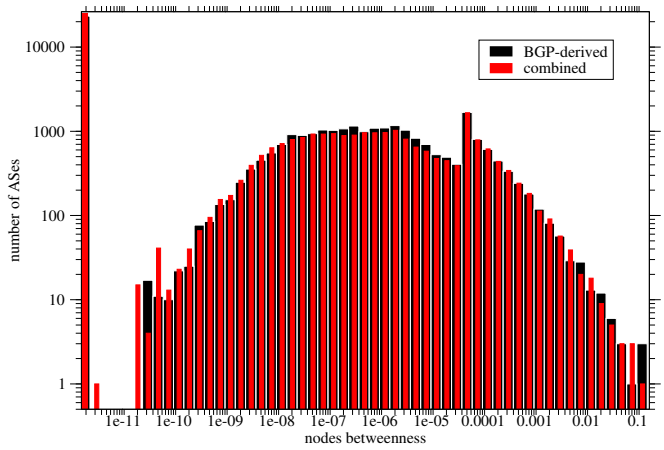


(a) The AS betweenness distributions are similar for both graphs. The leftmost column includes ASes with betweenness of 0.

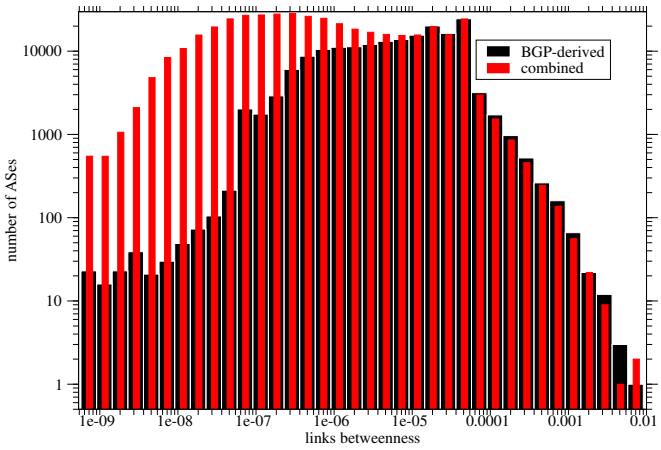


(b) The link betweenness distribution. Most new links added in the combined graph have low betweenness.

**Figure 6: AS and link betweenness, i.e., fraction of shortest paths that include the AS or link, for December 2014 data set (Section 3.2.2).**

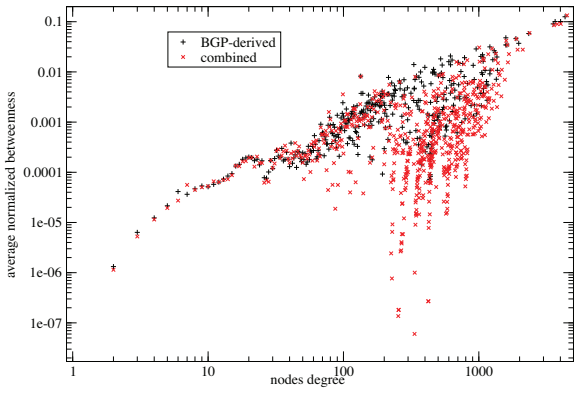


**Figure 7: The average AS betweenness as a function of AS degree. The betweenness of most ASes did not change. For some nodes, their degrees increased, but their betweenness remained the same, resulting in apparently lower betweenness values for a subset of nodes in the range of degrees between 100 and 1000 in the combined graph.**

*for transit*: $A$'s customer cone contains the AS $A$ itself, its customers, its customers' customers, and so on. More rigorously, we define $AS_A$'s customer cone to include $AS_A$ and the set of customers where $AS_{customer}$ is a customer of $AS_A$ or recursively a customer of $AS_A$'s customer, and $AS_A$ announces $AS_{customer}$ to a peer or provider of $AS_A$ as observed in BGP paths [3].

$AS_A$'s ***peer cone*** is the *set of ASes it can reach through its customer links or its peers' customer links*; that is, the peer cone includes the set of ASes in $AS_A$'s customer cone and in its peers' customer cones. In economic terms, $AS_A$'s peer cone is the set of ASes that $AS_A$ can reach without paying its provider. Note that this definition does not include the "observed in BGP paths" condition, since we have many peer links in the combined graph without corresponding observed BGP paths.

The customer cone by definition is constrained by observed paths, and the combined graph only includes additional AS links, but no new paths, so the customer cones are exactly the same in both graphs (Figure 8(a)). The customer cones tend to be rather small: 99% of ASes have a customer cone size less than 46. The peer cone sizes (Figure 8(b)) are distributed much more widely, with 10% of ASes having cone sizes greater than 738 in the BGP graph and greater than 918 in the combined graph. In both graphs, fewer than 1% of ASes have peer cone size greater than $17,000$.

For both kinds of cones, the distribution of cone sizes is dominated by the maximum at 1: 85% of ASes have only one AS (itself) in its customer cone and 78% of ASes have only itself in its peer cone (Figure 8). These ASes constitute the *edge* of the graph. Additional links in the combined graph caused only a miniscule change in the edge ASes: only 0.4% of them increased the size of their peer cones.

## 4.2 AS Ranking

Internet full connectivity critically depends on a small clique of large Internet Service Providers (ISPs) (represented as ASes in the AS topology) who base routing decisions on a complete table of Internet routes, and do not use anybody for transit (*Full Table No Transit (FTNT)* ASes). Such ISPs can reach every AS in the topology either through one of their customers or one of their peers' customers. Determining the members of this clique, which essentially forms the Internet core, is the first step in our AS relationship inference algorithm [3]. Discussions with operators have consistently validated the clique membership.

For the December 2014 data used in this report, we used the methodology described in [3] to estimate that there were 16 FTNT ASes. Specifically, we found the largest clique among the top 50 ASes ranked by transit degree (defined as the number of ASes for which the given AS was observed providing transit in BGP paths). These ASes matched the list of the "Tier 1" ASes by Wikipedia in 2014 [15]. Although larger ISPs often use more then one AS, the bulk of their connectivity is usually captured reasonably well by their primary AS – the one with the largest customer cone. Therefore, we consider the primary ASes as reasonable proxies for the corresponding FTNT ISPs.

Some metrics discussed in Sections 3.1 and 3.2 rank ASes as ei-

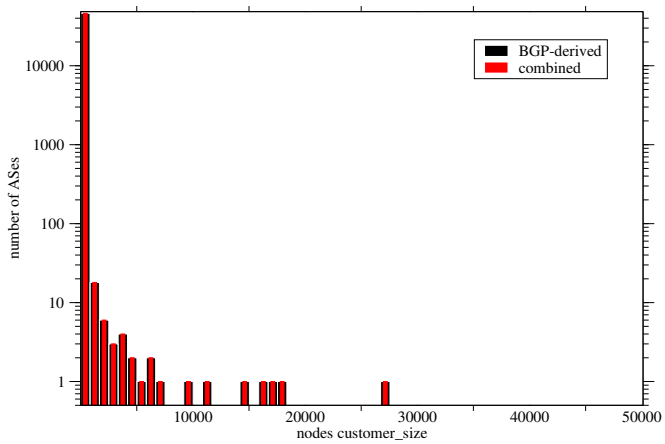


(a) Customer cone size distribution. Customer cones are calculated based on observed paths, and the combined graph does not entail any new paths, so the customer cones are the same in each graph.

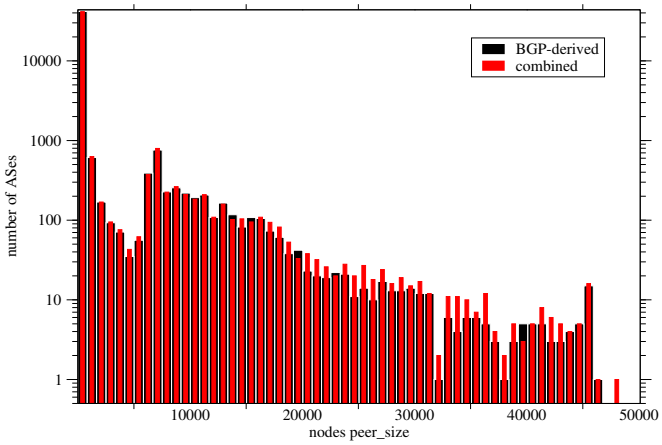


(b) The peer cone size distributions are similar for sizes $< 15,000$. The number of ASes with peer cone sizes $> 15,000$ noticeably increased in the combined graph.

**Figure 8: Customer and peer cone size distributions for December 2014 data set (Section 4.1).**

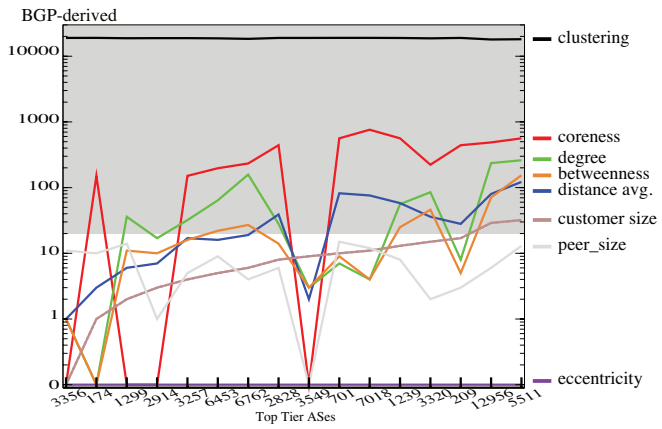


(a) BGP-derived graph

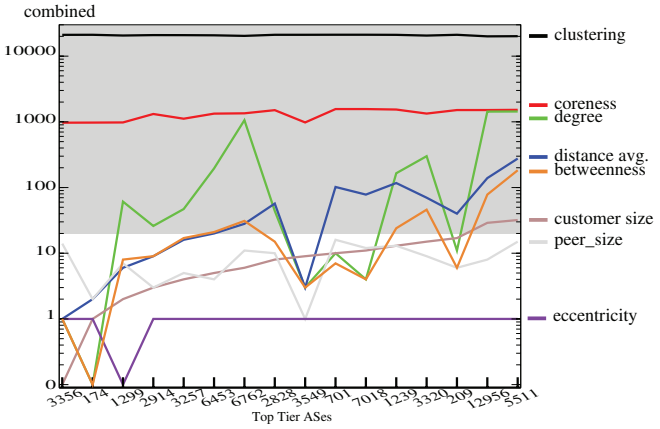


(b) combined graph

**Figure 9: Ranking of the Full Table No Transit (FTNT) ASes by the different metrics. The gray area in each graph shows ranking above 20. The eccentricity metric does not differentiate nodes, since 346 ASes had an eccentricity value of 6. Peer cone size would appear to be a better of AS routing influence then customer cone, with two additional FTNT ASes ranked in the top 20. Coreness and Clustering are the furthest from being able to capture a metric of AS importance to the global routing system.**

ther larger or more central. For example, transit degree and customer/peer cone sizes characterize the relative size of ASes, in terms of how much of the Internet address space they serve as providers or peers. Coreness, betweenness, average distance, eccentricity, and clustering indicate the relative centrality or importance of a given AS. Would these metrics rank highly those ASes known to be at the top of the Internet hierarchy in terms of customer/provider relationships? In other words, how well do the graph-theoretical metrics match and approximate our understanding of the Internet ecosystem?

| | BGP derived | combined |
|---|---|---|
| **number nodes** | | |
| total | 46,177 | 46,321 |
| **number links** | | |
| customer/provider | 93,539 | 93,539 |
| peer | 83,852 | 325,312 |
| total | 177,391 | 418,851 |
| **AS degree** (number of neighbors) | | |
| average | 7.7 | 18.1 |
| maximum | 4333 | 4484 |
| **average neighbor degree** | | |
| average (ave.) | 692.8 | 766.4 |
| ave. (normalized) | 0.015 | 0.02 |
| max. (normalized) | 0.023 | 0.024 |
| **clustering** (number of triangles) | | |
| mean local clustering | 0.28 | 0.34 |
| clustering coefficient | 0.05 | 0.47 |
| **coreness** | | |
| average | 2.96 | 11.93 |
| maximum | 76 | 409 |
| minumum | 0 | 0 |
| **distance** (number of ASes) | | |
| average | 3.75 | 3.68 |
| standard deviation | 0.85 | 0.83 |
| **eccentricity** (number of ASes) | | |
| average | 7.61 | 7.31 |
| diameter | 11 | 10 |
| radius | 6 | 5 |
| **node betweenness** (normalized) | | |
| average | $6.0\text{x}10^{-5}$ | $5.8\text{x}10^{-5}$ |
| minimum | 0 | 0 |
| maximum | 0.126 | 0.134 |
| **link betweenness** (normalized) | | |
| average | $21\text{x}10^{-6}$ | $9\text{x}10^{-6}$ |
| minimum | $9.4\text{x}10^{-10}$ | $9.3\text{x}10^{-10}$ |
| maximum | 0.0084 | 0.0071 |
| **customer cone size** (number of ASes) | | |
| minimum | 1 | 1 |
| maximum | 26,870 | 26,870 |
| average | 7.6 | 7.6 |
| percentage with size 1 | 84.8% | 84.8% |
| 90th percentile | 2 | 2 |
| 99th percentile | 46 | 46 |
| **peer cone size** (number of ASes) | | |
| minimum | 1 | 1 |
| maximum | 46,403 | 48,203 |
| average | 937 | 1036 |
| percentage with size 1 | 77.5% | 77.1% |
| 90th percentile | 738 | 918 |
| 99th percentile | 17,219 | 19,257 |

**Table 2: Summary statistics of both graphs.**

| | correlation between BGP-derived and combined |
|---|---|
| **eccentricity** | 0.7229 |
| **degree** | 0.7481 |
| **coreness** | 0.7715 |
| **clustering** | 0.8781 |
| **peer size** | 0.9660 |
| **betweenness** | 0.9666 |
| **distance avg.** | 0.9846 |
| **customer size** | 1.0000 |

**Table 3: Correlation between AS ranking in the BGP-derived and combined graphs. Rankings by eccentricity are the most different. Rankings by peer cone size, betweenness, and average distance are rather similar.**

We define an AS's rank for a given metric as the number of ASes with a value of this metric larger or more central than itself. This definition means that the rank value of the "top" AS for a given metric is 0. We rank cone size, degree, coreness, betweenness, and clustering in decreasing order, average distance and eccentricity in increasing order. Figure 9 shows how the metrics we have discussed rank the FTNT set of ASes. The $X$-axis plots ASes by decreasing size of their customer cones. Thus, the ranking by this metric (the brown line) grows monotonically. However, the ASes forming the clique do not necessarily have the largest customer cone sizes: only 13 FTNT ASes are ranked in the top 16 ASes by customer cone size; the last three (Qwest 209, Telefonica 12956 and Opentransit 5511) are ranked 17, 29, and 32 (in both graphs, since customer cones did not change).

Ranking by peer cone size (the grey line in Figure 9) most closely matches the FTNT clique: all clique member ASes are ranked 15 or higher. This tight relationship is unsurprising given the relationship between transit degree and inference of peer and customer relationships. In both graphs, the largest clique formed in the top 50 ranked ASes includes the ASes from the transit degree clique.

Ranking by betweenness (the orange line in Figure 9) is the third best at reflecting our notion of AS importance: in both graphs, nine of the 16 FTNT ASes are in the top 16 ranked by this metric, and only one AS is ranked below 100.

Surprisingly, the FTNT ASes are not among the ASes with the highest degree: this metric (the green line) places only six FTNT ASes (the same in both graphs) in the top 16, the rest ranked between 17 and 261 in the BGP graph, and between 26 and 1435 in the combined graph.

Coreness (the red line in Figure 9) yields a poor representation of the Internet routing hierarchy. Although four FTNT ASes have the highest value of coreness in the BGP-based graph (and corresponding rank of 0), this metric ranks the other 12 clique members far below 100. It performs even worse for the combined graph, ranking all 16 FTNT ASes near 1000 or lower. This behavior derives from the fact that most links in the combined graph connect medium ASes, which increases their coreness, and decreases the relative coreness of other nodes in the combined graph, ranking all 16 FTNT ASes at around 1000 and lower.

Distance-based metrics, i.e., the average distance from a given

| | BGP-derived | | | | | | | | | | | | | | | | | combined | | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | median | 1299 | 3356 | 3549 | 2914 | 3257 | 174 | 6453 | 3320 | 6762 | 12956 | 209 | 2828 | 701 | 1239 | 7018 | 5511 | median | 1299 | 3356 | 3549 | 2914 | 3257 | 174 | 6453 | 3320 | 6762 | 12956 | 209 | 2828 | 701 | 1239 | 7018 | 5511 |
| eccentricity | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 |
| customer size | 8 | 2 | 0 | 9 | 3 | 4 | 1 | 5 | 15 | 6 | 29 | 17 | 8 | 10 | 13 | 11 | 32 | 8 | 2 | 0 | 9 | 3 | 4 | 1 | 5 | 15 | 6 | 29 | 17 | 8 | 10 | 13 | 11 | 32 |
| peer size | 6 | 14 | 11 | 0 | 1 | 5 | 10 | 9 | 2 | 4 | 6 | 3 | 6 | 15 | 8 | 12 | 13 | 8 | 7 | 14 | 1 | 3 | 5 | 2 | 4 | 9 | 11 | 8 | 6 | 10 | 16 | 13 | 12 | 15 |
| betweenness | 11 | 11 | 1 | 3 | 10 | 16 | 0 | 22 | 46 | 27 | 71 | 5 | 14 | 9 | 25 | 4 | + | 9 | 8 | 1 | 3 | 9 | 17 | 0 | 21 | 46 | 31 | 78 | 6 | 15 | 7 | 24 | 4 | + |
| distance avg | 19 | 6 | 1 | 2 | 7 | 17 | 3 | 16 | 36 | 19 | 80 | 28 | 39 | 82 | 58 | 76 | + | 28 | 6 | 1 | 3 | 9 | 16 | 2 | 20 | 70 | 28 | + | 40 | 57 | + | + | 78 | . |
| degree | 28 | 36 | 1 | 3 | 17 | 32 | 0 | 64 | 85 | + | . | 8 | 28 | 7 | 56 | 4 | . | 44 | 61 | 1 | 3 | 26 | 47 | 0 | + | . | . | . | 11 | 44 | 10 | + | 4 | . |
| coreness | . | 0 | 0 | 0 | 0 | + | + | + | . | . | . | . | . | . | . | . | . | . | . | . | . | . | . | . | . | . | . | . | . | . | . | . | . | . |
| clustering | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |

rank key: 0-19 $^{0-19}$ 20-99 $^{20-99}$ + $^{99-199}$ . $^{201-2000}$ $^{2001+}$

**Table 4: AS Ranking. Lower ranked ASes are intuitively more central to the graph, based on the different metrics of the Full Table No Transit (FTNT) ASes. Light gray cells have ranks between 0 and 19. Ranks with values between 0 and 99 are show as numbers. Ranks between 100 and 199 are shown with a '+'. Ranks between 200 and 1999 have a '.'. Higher ranks are left blank. A metric that captures the economic independence of an AS would ideally rank all FTNT ASes in the top 20. Although eccentricity appears to have the best ranking, ranking all FTNT ASes below 20, an additional 346 ASes also have the same value of eccentricity, rendering it a useless differentiator. Peer cone size appears a more economically relevant metric then customer cone, with two additional FTNT ASes ranked in the top 20. Coreness and Cluster are the worst; the clustering metric does not rank a single AS in the top 20.**

| | set of ASes with | | | |
|---|---|---|---|---|
| **metric** | **peripheral** | | **central** | |
| ecccentricity | **periphery** | highest | **center** | lowest |
| betweenness | **outskirt** | low | **interior** | high |
| coreness | **fringe** | lowest | **core** | highest |
| cone size | **edge** | 1 | **clique** | member* |

*member of largest clique in top 100 by transit degree

**Table 5: Labels and rank order for different metrics use to define peripheral and central regions.**

AS to all other ASes in the graph, and eccentricity, also do not match our notion of an AS' economic importance. Distance places only five of the FTNT ASes in the top 16. In the BGP-derived graph, all FTNT ASes (and 330 other ASes!) have the same (highest) eccentricity value of 6, and thus all are ranked 0 by this metric. In the combined graph, the eccentricity of one AS (1299 – Telia) became 5. (Telia is the largest European AS and the European IXPs data contributed more than 50% of the new links to the combined graph – see Section 2.) Therefore, in Figure 9(b) only this AS has rank 0, while all other ASes have rank 1. We conclude that this metric does not provide a particularly informative ranking.

Clustering is a poor measure of AS centrality/importance. FTNT ASes have rather low clustering: this metric ranks all of them lower than $10,000$ (the black line in Figure 9). Small clustering values are unsurprising because these ASes at the top of the Internet hierarchy have high degrees but their immediate neighbors tend to be smaller than they are, and generally do not have the resources to interconnect with each other. In contrast, the medium ASes, who both use and provide transit, have smaller degrees, but are more likely to connect to other medium-sized ASes as they have a strong economic incentive to peer with each other and reduce the amount of traffic they must pay to send via the FTNT ASes (Figure 3(b)).

To elucidate the differences in rankings between the two graphs, we use the Pearson product-moment correlation coefficient [16]: perfectly correlated rankings would have a coefficient of 1, and rankings with no correlation would have a value of 0. Table 3 shows that the largest differences are in rankings by eccentricity and degree, with correlation values of 0.72 and 0.75, correspondingly.

## 4.3 Peripheral and central regions of the graphs

In an AS graph annotated with AS business relationships, ASes with a customer cone size of 1 form the edge of the graph, while the clique members, which we infer as the FTNT ASes, constitute the most central part of the graph. Other metrics, such as coreness, eccentricity, and node betweenness also distinguish the peripheral and central parts of a graph.

Coreness-based definitions are the *fringe* and the *core*: the sets of ASes with the lowest and the highest coreness. The core of the BGP graph consisted of 151 ASes with a coreness of 76, while in the combined graph the core grew to 422 ASes, all of whom had a coreness of 409. At the same time, the fringe remained practically unaffected, shrinking from 15,533 ASes in the BGP graph to 15,514 ASes in the combined graph (a 0.1% decrease).

Eccentricity defines the *periphery* as the set of ASes with the highest eccentricity (equal, by definition, to the graph diameter), and the *center* as the set of ASes with the lowest eccentricity (the graph radius). In the BGP graph, the center consisted of 346 ASes with eccentricity of 6. Additional links in the combined graph increased the number of ASes with an eccentricity of 6 from 346 to 1286. However, the eccentricity of just one AS, 1299 Telia, dropped to 5, making this sole AS the formal center of the combined graph. The graph's diameter shrank from 11 to 10, but the graph periphery expanded from 57 ASes in the BGP graph to 93 ASes in the combined one.

There is no rigorous betweenness-based definition of periphery vs. center, but the shape of the node betweenness distribution (Figure 6(a)) suggests peripheral and central parts of the graph based on this metric. We call those the *outskirt* and *interior*. 50% of ASes in the BGP-based graph and 54% of ASes in the combined graph have betweenness of 0 (the leftmost peak in Figure 6(a)) and comprise the graph outskirt. The shape of the distribution changes for betweenness values $> 4 \times 10^{-5}$, where the distribution approximates a power law function with exponent of 0.88. This region includes 8.2% of ASes in the BGP-based graph and 10.0% of ASes in the combined graph, representing the graph interior.

Table 6 summarizes the different views of the macroscopic structure of the two graphs. We conclude that none of the graph-theoretic metrics we analyzed provides a meaningful approximation of the peripheral and central graph structures that emerge from analysis of AS business relationships.

| | | number of ASes (percentage of the total) | | | |
|---|---|---|---|---|---|
| | **metric** | **customer cone** | **coreness** | **eccentricity** | **node betweenness** |
| Peripheral part of the graph | BGP-derived | 39,142 (85%) | 15,533 (34%) | 57 (0.12%) | 23,343 (51%) |
| | combined | 39,142 (85%) | 15,514 (33%) | 93 (0.20%) | 25,053 (54%) |
| Central part of the graph | BGP-derived | 16 (0.035%) | 151 (0.32%) | 346 (0.75%) | 4639 (9.99%) |
| | combined | 16 (0.035%) | 422 (0.91%) | 1 (0%) | 4637 (10.04%) |

**Table 6: The peripheral and central parts of AS graphs as defined by different metrics. Depending on the metric, the results are wildly scattered. None of the graph-theoretical metrics provides a good approximation to the partitioning based on the actual AS relationships.**

# 5. CONCLUSIONS

The limitations of a BGP-inferred AS-level topology are well understood, but the impact of various types of missing links on various topology properties and characteristics remains an open question. We examined how the introduction of 241,459 additional peering links (a 136% increase) and 144 additional nodes (0.3% increase) inferred from traceroute and Internet eXchange (IX) data changed the topological properties of the AS-level graph originally derived from only BGP data.

Only 6.8% of the ASes in the original BGP-based graph gained additional links. Of those ASes, links were primarily added to medium degree ASes. For example, ASes with degrees between 20 and 200 make up only 4.5% of ASes in the original BGP-based graph, but represent 44.2% of ASes that got at least one additional link. The addition of new links increased the average degree from 7.6 to 18.1 and the fraction of ASes with degrees $> 230$ grew from 0.5% to 2.9% of the graph.

Given that the new links predominantly came from medium degree nodes, it is unsurprising that the number of ASes classified as edge ASes remained largely unaffected. For all four metrics we used to define the peripheral part of the graphs (customer cone, coreness, eccentricity, node betweenness), the change of the relative size of this part was 3% or less between the BGP-based and the combined graphs (see Table 6).

The relative ranking of the ASes was also largely unchanged across most of the metrics. For customer cone size, average distance, betweenness, and peer cone size, correlation between the ranking in the BGP-derived and the combined graphs exceeds 0.97. Only for degree and eccentricity is the correlation between rankings less than 0.75. The Full Table No Transit (FTNT) ASes were only ranked in the top 30 by cone sizes and eccentricity. Only clustering and coreness rankings never had at least 9 FTNT's in their top-ranked ASes.

If we had to distill these insights into a single sentence, it would be that for the very large degree and very small degree ASes, BGP measurements tend to capture their connectivity well, but for many middle-degree ASes, additional connectivity is revealed by traceroute-based inferences that is not visible in global BGP data repositories.

# 6. ACKNOWLEDGEMENT

The work was funded by the Department of Homeland Security (DHS) Science and Technology Directorate, Cyber Security Division (DHS S&T/CSD) Broad Agency Announcement 11-02 and SPAWAR Systems Center Pacific via contract number N66001-12-C-0130, and by Defence Research and Development Canada (DRDC) pursuant to an Agreement between the U.S. and Canadian governments for Cooperation in Science and Technology for Critical Infrastructure Protection and Border Security. The work represents the position of the authors and not necessarily that of DHS or DRDC.

# 7. REFERENCES

[1] "University of Oregon Route Views Project." `http://www.routeviews.org/`.

[2] "RIPE NCC RIS." `https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris`.

[3] M. Luckie, B. Huffaker, k. claffy, A. Dhamdhere, and V. Giotsas, "AS Relationships, Customer Cones, and Validation," in *Internet Measurement Conference (IMC)*, Oct 2013.

[4] "Archipelago (Ark) Measurement Infrastructure." `http://www.caida.org/projects/ark`.

[5] V. Giotsas, M. Luckie, B. Huffaker, and k. claffy, "Inferring Complex AS Relationships," in *Internet Measurement Conference (IMC)*, Nov 2014.

[6] M. Luckie, A. Dhamdhere, B. Huffaker, D. Clark, and k. claffy, "bdrmap: Inference of Borders Between IP Networks," in *Internet Measurement Conference (IMC)*, 2016.

[7] K. Chen, D. Choffnes, R. Potharaju, Y. Chen, F. Bustamante, D. Pei, and Y. Zhao, "Where the sidewalk ends: extending the internet as graph using traceroutes from P2P users," in *CoNEXT*, 2009.

[8] "Peering db." `www.peeringdb.com`.

[9] A. Lakhina, J. W. Byers, M. Crovella, and P. Xie, "Sampling biases in IP topology measurements," in *INFOCOM*, Apr. 2003.

[10] V. Giotsas, S. Zhou, M. Luckie, and k. claffy, "Inferring Multilateral Peering," in *ACM SIGCOMM Conference on emerging Networking EXperiments and Technologies (CoNEXT)*, Aug 2013.

[11] R. Oliveira, D. Pei, W. Willinger, B. Zhang, and L. Zhang, "The (in)completeness of the observed internet aslevel structure," *IEEE/ACM Transactions on Networking*, 2010.

[12] P. Mahadevan, D. Krioukov, M. Fomenkov, B. Huffaker, X. Dimitropoulos, k. claffy, and A. Vahdat, "The Internet AS-Level Topology: Three Data Sources and One Definitive Metric," *ACM SIGCOMM Computer Communication Review*, vol. 36, Jan 2006.

[13] "topostats." `http://www.caida.org/tools/utilities/topostats/`.

[14] U. Brandes, "A faster algorithm for betweenness centrality," *The Journal of Mathematical Sociology*, vol. 25, Dec 2000.

[15] "Tier 1 network." `https://en.wikipedia.org/wiki/Tier_1_network`.

[16] "Pearson product-moment correlation coefficient." `https://en.wikipedia.org/wiki/Pearson_product-moment_correlation_coefficient/`.