Internet Topology

Yihua He
yhe@cs.ucr.edu

Georgos Siganos
siganos@cs.ucr.edu

Michalis Faloutsos
michalis@cs.ucr.edu

University of California, Riverside

Contents

Glossary

Autonomous System (AS): An Autonomous System is a connected group of one or more IP prefixes run by one or more network operators which has a single and clearly defined routing policy. A unique AS number (or ASN) is allocated to each AS for identification purpose in inter-domain routing among ASes. For example, an organization, such as an ISP or a university, is an example of an AS. Some organization may have more than one ASes and thus have more than one AS numbers.

BGP (Border Gateway Protocol): The Border Gateway Protocol (BGP) is the *de facto* routing protocol used in the Internet to exchange reachability information among ASes and interconnect them. The current BGP is version 4.

Degree: The degree of an AS (or a node) is the number of neighbors of this AS (or node).

CCDF: The Complementary Cumulative Distribution Function (CCDF) of a degree, is the percentage of nodes that have degree greater than the degree.

Degree Rank: The degree rank of a node is its index in the order of decreasing degree.

Eigenvalue: Let A be an $N \times N$ matrix. If there is a vector $X \in \mathcal{R}^N \neq 0$ such that $AX = \lambda X$ for some scalar λ, then λ is called the eigenvalue of A with corresponding eigenvector X.

1 Definition

The *Internet Topology* is the structure of how hosts, routers or Autonomous Systems are connected to each other. Majority of the existing Internet topology research focuses on AS-level. This is because: (1) AS-level Internet topology is at the highest granularity of the Internet. Other levels of Internet topology partially depend on AS-level topology. (2)The AS-level Internet topology is relatively easy to obtain. Other levels of topology are sometimes regard as private

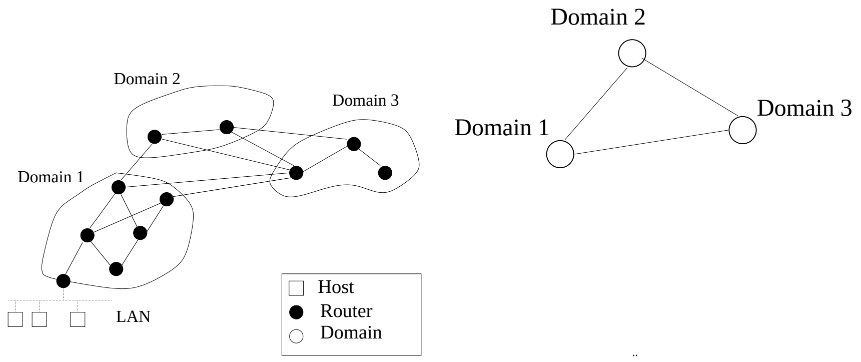

(a) Router level

(b) AS level

Figure 1: The structure of Internet at two levels

information and they are harder to get. (3)AS-level topology is not directly engineered by human; instead, it is driven by technological and economical forces.

The research on Internet topology is driven by the explosive growth of the Internet, which has been accompanied by a wide range of inter-networking problems related to routing, resource reservation and administration. The study of algorithms and policies to address such problems often requires topological information and models. In 1999, Faloutsos brothers [1] discovered that the seemingly random Internet topology does follow some rules: it follows power-law distributions. This finding revolutionized the research on Internet topology and generated much follow-up works.

2 Introduction

The Internet can be decomposed into connected subnetworks that are under separate administrative authorities, as shown in Figure 1. These subnetworks are called *domains* or *Autonomous Systems* (ASes). The Internet community develops and employs different routing protocols inside an AS and between ASes. An intra-domain protocol, such as RIP, IS-IS, or OSPF, is limited within an AS, while an inter-domain protocol, such as BGP, runs between ASes. This way, the topology of the Internet can be studied at two different granularities. At the **router level**, each router is represented by a node [2], and a direct connection (either inter-domain or intra-domain) between any pair of routers is represented by an edge. At the AS level, each AS is represented by a single node [3] and each edge is an inter-domain interconnection. The study of the topology at both levels is equally important.

This article majorly focus on the AS level Internet topology. Note that, at this level, only one edge exists between two nodes, although there may be multiple connections between two ASes in practice. This limitation is majorly due to the nature of the data that the community has.

Network	Next Hop	(Other Fields)	Path
2.0.0.0/24	157.130.10.233	(...)	701 1299 34211 41856 41856

Table 1: An excerpt of an entry in typical Cisco "sh ip bgp" format BGP table dumps from BGP collector "route-views.oregon-ix.net" on May 1, 2007.

There are multiple benefits from understanding the topology of the Internet. This is motivated by questions like the following "What does the Internet look like?" "Are there any topological properties that don't change in time?" "How will it look like a year from now?" "How can I generate Internet-like graphs for my simulations?". Modeling the Internet topology is an important open problem despite the attention it has attracted recently. Paxson and Floyd consider this problem as a major reason why we don't know how to simulate the Internet [4]. An accurate topological model can have significant impact on network research. First, we can design more efficient protocols that take advantage of its topological properties. Second, we can create more accurate artificial models for simulation purposes. And third, we can derive estimates for topological parameters that are useful for the analysis of protocols and for speculations of the Internet topology in the future.

In the rest of this article, we will first review how and where the Internet topology information is collected in Section 3. We also compare the pros and cons of different data sources and investigate the reason why they have such difference. In Section 4, we will review the power-laws of the Internet, which is one of the most important discoveries of the Internet topology. The power-laws lead to significant follow-up works in modeling the Internet topology, and we will discuss these models in Section 5 and Section 6. Then we exhibit the techniques to discover the complete Internet topology in Section 7. At last, we discuss future research directions of the Internet topology in Section 8.

3 Data Sources and Comparison

We first describe the data sources for collecting AS-level Internet topology. All of these sources and methods have their shortcomings, which will be compared at the end of this section.

3.1 BGP routing tables

BGP routing table dumps are probably the most widely used resource that provides information on the AS Internet topology. Normally, these routing table dumps are obtained from special BGP collectors[1], each of which connects with one or more Internet backbone routers in different ASes with special agreements.

[1] Normal BGP routing software runs on the BGP collectors. Therefore, each BGP collector has all functionalities of a normal BGP router; routing table dumps can be obtained from these BGP collectors.

These BGP collectors do not advertise any prefix (i.e., IP blocks) to the Internet, while they are configured to receive all routes that the Internet backbone routers advertise to them. Therefore, these collectors are totally passive, and they have no effect to the global Internet. Periodically, each collector dumps its full routing tables to Internet archives, which are available for download.

Table 1 shows an entry of typical Cisco “sh ip bgp” format BGP table dumps from BGP collector “route-views.oregon-ix.net” on May 1, 2007. This entry indicates that, the destination network 2.0.0.0/24 can be reached via AS path “701 1299 34211 41856”. Therefore, the instance of Internet topology should include four ASes (AS701, AS1299, AS34211 and AS41856) and three links (AS701-AS1299, AS1299-AS34211 and AS34211-AS41856). Typically a BGP collector’s routing table dump has more than a hundred thousands such entries from each peer AS; the total number of entries often exceeds several millions. The number of ASes or routers that each BGP collector varies from a few to approximately one hundred. There are two well-maintained BGP routing table collector agent: Oregon Routeviews[5] and RIPE RIS[6]. Each of these agents maintains a number of BGP collectors over the world.

Besides routing tables, a BGP collector may also dump routing updates received from its peers periodically. Routing updates have a similar format to routing tables. A BGP update message displays the current route to a prefix, and therefore, a collection of BGP updates is able to reveal the dynamic of BGP routing.

3.2 Traceroute

Traceroute[7] is a handy debugging program to discover the route that IP datagrams follow from one host to another. Traceroute takes advantage the fact that each router has to decrement the TTL (Time To Live) field by 1 for each IP packet pass through it, and each router has to discard any IP packets with TTL=0 and send an ICMP “time-exceeded” error message back to the sender of the original IP packet. The original purpose of the action is to prevent IP packets from circulating the network for ever. The operation of traceroute take advantage of this feature: it sends TTL increasing IP packets to the destination. These packets will expire at the routers along the path it reach the destination. Since each router along the path will send ICMP “time-exceeded” message back to the traceroute source, the identities of these routers (or more precisely, the outgoing IP interfaces of those routers) can be discovered. Although there is no guarantees that two consecutive IP packets will traverse the same route to the same destination, most often they do.

Skitter[8] is a part of CAIDA’s [9] topology measurement project. CAIDA [9] maintains a set of (about 20) active monitors distributed around the globe. Each monitor uses a modified version of traceroute to probe a large set of IP addresses which nearly cover the whole IP address space. Rocketfuel [10] is a topology discovery project from University of Washington. Rocketfuel use a larger number (a few hundreds) of traceroute sources from public traceroute servers as their sources. Therefore Rocketfuel has significant more number of vantage points. However, due to restrictions of the public traceroute servers, the rate of traceroute probing is limited in Rocketfuel. Therefore Rocketfuel is better at probing specific ISP networks but not the whole Internet. Another promising project is NetDimes [11] from Telaviv University. To increase the number of vantage points, NetDimes distribute a large number (tens of thousands) of probing agents to global Internet users on volunteer basis. These agents perform traceroutes according to the NetDimes center controls. Since the agents are most volunteers, coordination is still hard in order to probe the Internet topology from anywhere at any given time.

All traceroute probes only reflect router-level topology. In order to obtain AS level Internet topology, the probed IP addresses have to be mapped to the ASes that they belong to. The conventional way to map an IP address to its AS number is by looking up the IP block in the BGP routing tables with the longest prefix match. For example, if there is an IP address 2.0.0.18 and the longest prefix that it matches in the routing table is 2.0.0.0/24 (as shown in Table 1), then the announcing AS, which is the AS appeared at the end of the "Path" field, (AS41856 in Table 1) is the AS that the IP 2.0.0.18 should be mapped to. However, the accuracy of such conversion may suffer in certain situations. Mao et. al. [12][13] discussed them in details.

3.3 Internet Routing Registry (IRR)

The need for cooperation between Autonomous Systems is fulfilled today by the Internet Routing Registries (IRR) [6]. ASes use the Routing Policy Specification Language (RPSL) [7] [8] to describe their routing policy, and router configuration files can be produced from it. At present, there exist 55 registries, which form a global database to obtain a view of the global routing policy. Some of these registries are regional, like RIPE or APNIC, other registries describe the policies of an Autonomous System and its customers, for example, cable and wireless CW or LEVEL3. The main uses of the IRR registries are to provide an easy way for consistent configuration of filters, and a mean to facilitate the debugging of Internet routing problems. From the registered routing export and import policies, Internet topology can be extracted from IRR. For example, in Table 2, an excerpt of *aut-num* record for AS3303 in IRR is shown. From the registered import and export policy in this excerpt, the Internet topology should include three ASes (AS3303, AS701 and AS1239), and two edges (AS3303-AS701 and AS3303-AS1239).

3.4 Data source comparison

BGP table dumps, especially the one from Oregon Routeview project, are the most widely used source for Internet topology study. An advantage of the BGP routing tables is that their link information is considered reliable. If an AS link appears in a BGP routing table dump, it is almost certain that the link exists. However, limited number of vantage points makes it hard to discover a more complete view of the AS-level topology. A single BGP routing table

```
aut-num:   AS3303
as-name:   SWISSCOM
descr:     Swisscom Solutions Ltd
descr:     IP-Plus Internet Backbone
...        ...
import:    from AS701 action pref=700; accept ANY
export:    to AS701 announce AS-SWCMGLOBAL
import:    from AS1239 action pref=700; accept ANY
export:    to AS1239 announce AS-SWCMGLOBAL
...        ...
```

Table 2: An excerpt of IRR in plain text format for AS3303.

has the union of “shortest” or, more accurately, preferred paths with respect to this point of observation. As a result, such a collection will not see edges that are not on the preferred path for this point of observation. Several theoretical and experimental efforts explore the limitations of such measurements [14][15]. Worse, such incompleteness may be statistically biased based on the type of the links: peer-to-peer links are more likely to be missing from BGP routing table dumps than provider-customer links, due to the selective exporting rules of BGP. Typically, a peer-to-peer link can only be seen in a BGP routing table of these two peering ASes or their customers. Thus, given a peer-to-peer edge, unless a BGP collector peers with a customer of either AS incident to the edge, the edge can not be detected from the table dumps of the BGP collector. A recent work [16] discusses in depth this limitation. Thus, apart from being incomplete, the measured graph may not fairly represent the different types of links. Further more, BGP table dumps are likely to miss alternative and back-up paths. By definition, a router advertises only the best path to each destination, namely an IP prefix. Therefore, the back-up paths will not show up in distant ASes unless the primary link breaks. To address the problem, a recent effort suggests the need for actively probing backup links [17].

BGP updates are used in previous studies[18][19] as a source of topological information and they show that by collecting BGP updates over a period of time, more AS links are visible. This is because as the topology changes, BGP updates provide transient and ephemeral route information. However, if the window of observation is long, an advertised link may cease to exist [18] by the time that we construct a topology snapshot. In other words, BGP updates may provide a superimposition of a number of different snapshots that existed at some point in time. Note that BGP updates are collected at the same vantage points as the BGP tables in most collection sites. Naturally, topologies derived from BGP updates share the same statistical bias per link type as from BGP routing tables: peer-to-peer links are only to be advertised to the peering ASes and their customers. This further limits the additional information that BGP updates can provide currently. On the other hand, BGP updates could be useful in revealing ephemeral backup links over long period of observation, along with erroneous BGP updates, which are not visible in the Internet at large, unless the primary link breaks down. To tell the two apart, we need highly targeted probes. Recently, active BGP probing[17] has been proposed for identifying backup AS links. This is a promising approach that could complement our work and provide the needed capability for discovering more AS links.

By using traceroute, one can explore IP paths and then translate the IP addresses to AS numbers, thus obtaining AS paths. Similar to BGP tables, the traceroute path information is considered reliable, since it represents the path[2] that the packets actually traverse. On the other hand, a traceroute server explores the routing paths from its location towards the rest of the world, and thus, the collected data has the same limitations as BGP data in terms of completeness and link bias. One additional challenge with the traceroute data is the mapping of an IP path to an AS path. The problem is far from trivial, and it has been the focus of several recent efforts [12][13].

Internet Routing Registry (IRR)[20] is the union of a growing number of world-wide routing policy databases that use the Routing Policy Specification Language (RPSL). In principle, each AS should register routes to all its neighbors (that reflect the AS links between the AS and its neighbors) with this registry. IRR information is manually maintained and there is no stringent requirement for updating it. Therefore, without any processing, AS links derived from IRR are prone to human errors, could be outdated or incomplete. However, the up-to-date IRR entries provide a wealth of information that could not be obtained from any other source. A recent effort [21] shows that, with careful processing of the data, we can extract a non-trivial amount of correct and useful information.

4 Power-Laws of the Internet

The power-laws for Internet topology are first observed by Faloutsos brothers [1], and later summarized in [22]. In those two papers, the authors have shown that the Internet topology at the AS level can be described efficiently with power-laws. The elegance and simplicity of the power-laws provide a novel perspective into the seemingly uncontrolled Internet structure.

Power-laws are expressions of the form $y \propto x^a$, where a is a constant, x and y are the measures of interest, and $\propto$ stands for "proportional to". Pareto was among the first to introduce power-laws in 1896 [23]. He used power-laws to describe the distribution of income where there are few very rich people, but most of the people have a low income. Another classical law, the Zipf law [24], was introduced in 1949, for the frequencies of the English words and the population of cities. More recently, power-laws have been observed in communication networks. First, power-laws have been observed in traffic [25][26][27]. In addition, the topology of the World Wide Web [28, 29] can be described by power-laws.

[2] An exception is when the route changes while a path is being explored by a traceroute.

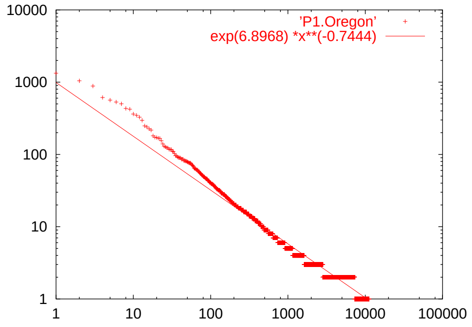

Figure 2: Log-log plot of the degree d_v versus the rank r_v in the sequence of decreasing degree.

Furthermore, power-laws describe the topology of peer-to-peer networks [30] and properties of multicast trees [31, 32, 33, 34].

The initial work on power-laws [1] has generated significant follow-up work. In fact, [1] is one of the top five most cited computer science papers published in 1999[35]. Various researchers verified the power-law observations with different datasets[36, 37, 38]. In addition, significant work has been devoted in understanding the origin [39], and generating power-law topologies [40, 39, 41, 42, 43, 44, 45, 46, 47].

For the Internet topology, three power-laws have been identified: a rank power-law, a degree power-law and an eigen power-law.

4.1 Rank power-law

Power-law 1 (rank exponent) *Given a graph, the degree, d_v, of a node v, is proportional to the rank of the node, r_v, to the power of a constant, $\mathcal{R}$:*

$$d_v \propto r_v^{\mathcal{R}}$$

Definition 1 *Let us sort the nodes of a graph in decreasing order of degree. We define the rank exponent, $\mathcal{R}$, to be the slope of the plot of the degrees of the nodes versus the rank the nodes in log-log scale.*

Figure 2 shows the (r_v, d_v) pairs in log-log scale after the nodes in an Internet topology are sorted in decreasing order of degree, d_v. The measured data is obtained from Oregon Routeviews [5] collector and represented by points, while the solid line represents the least-squares approximation. A striking observation is that the plots are approximated well by linear regression. The correlation coefficient is over 0.97 in this case. The authors of [22] also have inspected more than 1000 Internet topology instances and over six year span between 1997 and

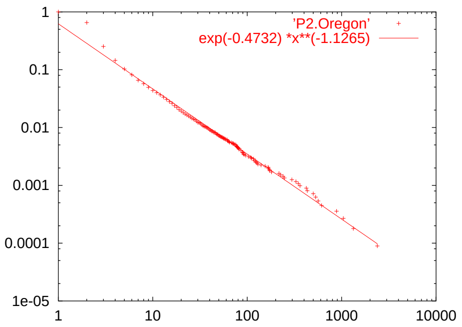

Figure 3: The log-log plot of D_d versus the degree for the Oregon topologies.

2003, and they found that for every instance of the inter-domain topology the correlation coefficient was always higher than 0.97. This linearity is unlikely to be a coincidence.

Intuitively, Power-law 1 correlate the degrees of the nodes and their rank and reflects a principle of the way domains connect. Such relationship can be used to calculate the number of edges as a function of the number of nodes for a given rank exponent. In fact, in a graph where Power-law 1 holds, the number of edges, E, of a graph can be estimated as a function of the number of nodes, N, and the rank exponent, R, as follows:

$$E = \frac{N}{2(R+1)}(1 - \frac{1}{N^{R+1}})$$

For additional discussion on estimates using this formula, see [1].

4.2 Degree power-law

Power-law 2 (degree exponent) *Given a graph, the CCDF, D_d, of an degree, d, is proportional to the degree the power of a constant, $\mathcal{D}$:*

$$D_d \propto d^{\mathcal{D}}$$

Definition 2 *We define the degree exponent, $\mathcal{D}$, to be the slope of the plot of the Cumulative degree of the degrees versus the degrees in log-log scale.*

In Figure 3, D_d is plotted versus the degree d in log-log scale. The major observation is that the plot is linear. The correlation coefficient is more than 0.996 for the data obtained from Oregon Routeviews[5]. Authors in [22] found that the degree power-law holds for all the instances they inspected from 1997 to 2003, with correlation coefficient higher than 0.99.

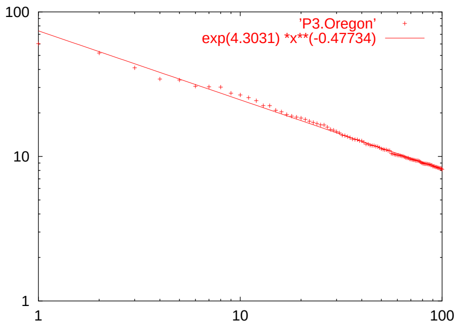

Figure 4: The eigenvalues plot for the Oregon topologies.

The intuition behind this power-law is that the distribution of the degree of Internet nodes is not arbitrary. The qualitative observation is that lower degrees are more frequent. The power-law manages to quantify this observation by a single number, the degree exponent. This way, the realism of a graph can be tested with a simple numerical comparison. If a graph does not follow Power-law 2, or if its degree exponent is considerably different from the real exponents, it probably does not represent a realistic topology.

The exponents of the rank and degree power-laws are shown to be related[48][49]. More specifically, in a perfect power-law distribution, the exponent of the rank power-law is equal to the multiplicative inverse of the exponent of the degree power-law. However, in reality, the two exponents hardly have such perfect relationship. The discrepancy could be attributed to measurement imperfections and inaccuracies. In that regard, both the rank and the degree power-laws characterize the degree distribution from different angles, and it is useful to report both exponents when characterizing a topology.

4.3 Eigen power-law

Power-law 3 (eigen exponent) *Given a graph, the eigenvalues, λ_i, are proportional to the order, i, to the power of a constant, $\mathcal{E}$:*

$$\lambda_i \propto i^{\mathcal{E}}$$

Definition 3 *We define the eigen exponent, $\mathcal{E}$, to be the slope of the plot of the sorted eigenvalues versus their order in log-log scale.*

Eigenvalues of a graph are the eigenvalues of its adjacency matrix. In Figure 4, the eigenvalues are plotted versus the their order in the decreasing sequence, in log-log scale. The eigenvalues are shown as points in the figure, and the solid lines are approximations using a least-squares fit. Similar observations with equally high correlation coefficients were observed for all instances obtained between 1997 and 2003 [22]. The plot is practically linear with a correlation coefficient of 0.996, which constitutes an empirical power-law of the Internet topology.

Eigenvalues are fundamental graph metrics. There is a rich literature that proves that the eigenvalues of a graph are closely related to many basic topological properties such as the diameter, the number of edges, the number of spanning trees, the number of connected components, and the number of walks of a certain length between vertices, as shown in [50]. Interestingly, Mihail et al. [51] show that there is a surprising relationship between the eigen exponent and the degree exponent: the eigen exponent is approximately the half of the degree exponent. In practice, the exponents obey adequately the mathematical relationship, although the match is naturally not perfect. All of the above suggest that the eigenvalues intimately relate to topological properties of graphs. However, it is not trivial to explore the nature and the implications of this power-law.

4.4 The doubts and settlement

There has been doubts and debate on whether the degree distribution of the Internet at the AS level follows a exact power-law [46][52]. The major concern is that by adding new edges discovered from a number of sources other than the most used Oregon Routeviews [5], the degree distribution of the Internet topology deviates from a perfect power-law.

There are at least two reasons for this debate. First, this debate is partly due to the absence of a definitive statistical test. In Figure 5 (a), the CCDF of node degrees is plotted for an Internet topology instance obtained from multiple resources, including Routeviews [5] and verified edges from IRR [20]. The distribution is highly skewed, and the correlation coefficient of a least square errors fitting is 98.9%. However, one could still use different statistical metrics and argue against the accuracy of the approximation [46]. Second, the answer could vary depending on which source we think is more complete and accurate, and the purpose or the required level of statistical confidence of a study. In Figure 5 (b), the CCDF is plotted for an Internet topology instance obtained from IRR after filtered by Nemecis [21]. The correlation coefficient is only 93.5%.

A recent paper [53] propose a reconciliatory divide-and conquer approach to explain and settle the debate: they propose to model separately the degree distribution according to the type of the edges: provider-customer and peer-to-peer[3]. In Figure 5, an indicative set of degree distribution plots are shown. The graphs obtained from multiple sources (Oregon Routeviews and traceroute

[3] A provider-customer edge means the two ASes incident to the edge have provider-customer relationship, i.e., one AS pays the other AS for traffic transit service. A peer-to-peer edge means the two AS incident to the edge have peer-to-peer relationship, i.e., these two ASes have mutual agreement that they carry traffic for each other with no or little fee. The classification of AS relationships can be inferred fairly accurately from the BGP routing tables by a number of algorithms[54][55][56][57]. Majority of the edges (approximately 80%) are provider-customer edges and most of the rest edges are peer-to-peer edges [53].

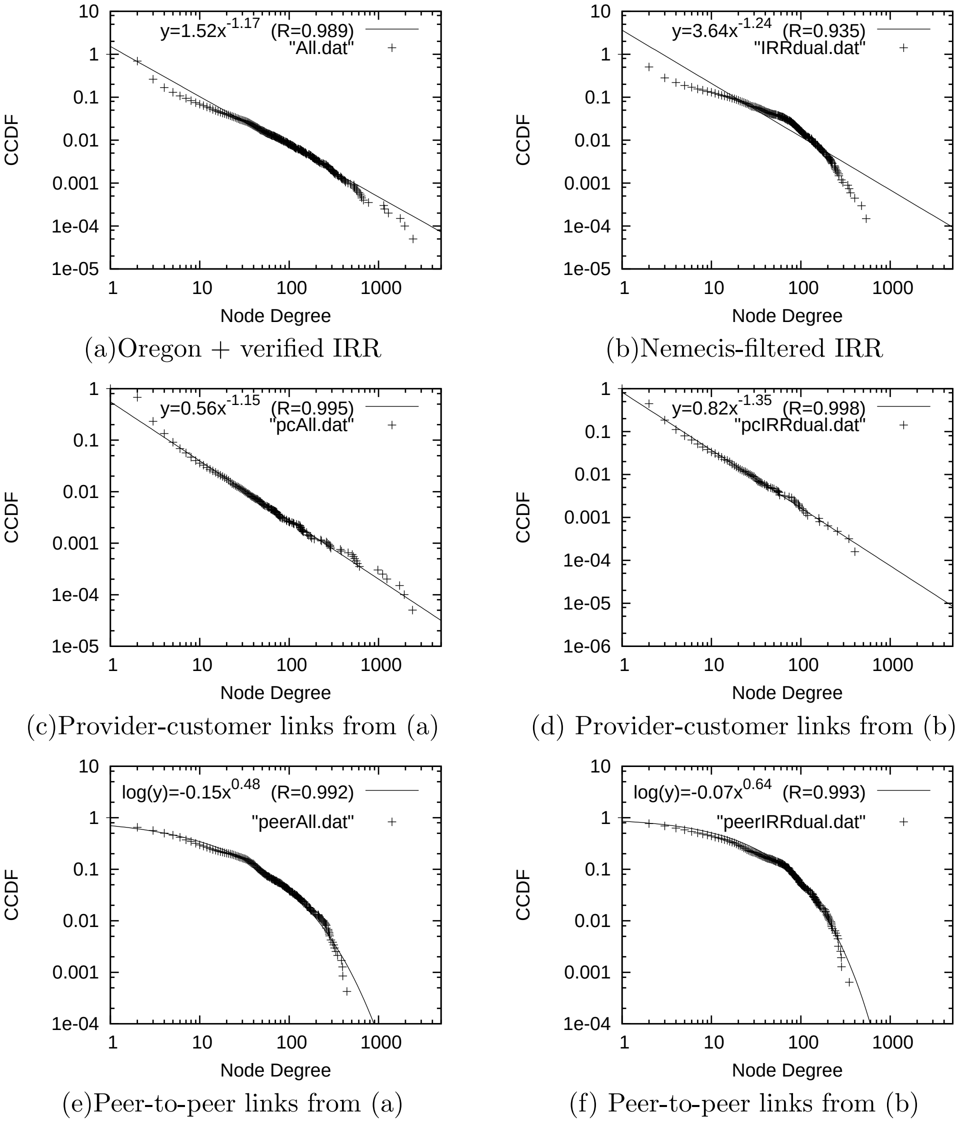

(a)Oregon + verified IRR

(b)Nemecis-filtered IRR

(c)Provider-customer links from (a)

(d) Provider-customer links from (b)

(e)Peer-to-peer links from (a)

(f) Peer-to-peer links from (b)

Figure 5: The degree distributions of *Oregon + verified IRR* (left) and *Nemecis-filtered IRR* (right) in the top row, their provider-customer degree distributions in the middle row, and their peer-to-peer degree distributions in the bottom row.

verified IRR links) are plotted in the left column ((a), (c), and (e)), and the topology obtained from Nemecis-filtered IRR are plotted in the right column ((b), (d), and (f)). The distributions for the whole graph are shown in the top row, the provide-customer edges only in the middle row, and the peer-to-peer edges only in the bottom row. The power-law approximation in the first two rows of plots and the Weibull approximation in the bottom row of plots are shown.

The following two properties can be observed from Figure 5: (1)The provider-customer-only degree distribution can be accurately approximated by a power-law. The correlation coefficient is 99.5% or higher in the plots of Figure 5 (d) and (e). Note that, although the combined degree distribution of the topology in IRR does not follow a power law as shown in Figure 5 (b), its provider-customer subgraph follows a strict power law in Figure 5 (d). (2)The peer-to-peer-only degree distribution can be accurately approximated by a Weibull distribution. The correlation coefficient is 99.2% or higher in the plots of Figure 5 (e) and (f). It is natural to ask why the two distributions differ. The following could be one of the explanations. Power-laws are related to the rich-get-richer behavior: low degree nodes "want" to connect to high degree nodes. For provider-customer edges, this makes sense: an AS wants to connect to a high-degree provider, since that provider would likely provide shorter paths to other ASes. This is less obviously true for peer-to-peer edges. If AS1 becomes a peer of AS2, AS1 does not benefit from the other peer-to-peer edges of AS2 due to routing policies[58]: an AS normally will not carry traffic from one of its peers to its other peers. Therefore, high peer-to-peer degree does not make a node more attractive as a peer-to-peer neighbor. The validity of this explanation is still under investigation[53].

4.5 Network analysis before power-laws

Before the discovery of power-law distribution of Internet topology, the metrics that had been used to describe graphs were mainly the node degree and the distances between nodes. Given a graph, the distance between two nodes is the number of edges along the shortest path between the two nodes. Most studies report minimum, maximum, and average values and plot the degree and distance distribution. Govindan and Reddy [3] study the growth of the inter-domain topology of the Internet between 1994 and 1995. The graph is sparse with 75% of the nodes having degrees less or equal to two. Pansiot and Grad [2] study the topology of the Internet in 1995 at the router level. The distances they report are approximately two times larger compared to those of Govindan and Reddy.

For graph generation purposes, Waxman introduced what seemed to be one of the most popular network models [59]. These graphs are created probabilistically considering the distance between nodes in a Euclidean sense. This model was successful in representing small early networks such as the ARPANET. As the size and the complexity of the network increased more detailed models were needed [60] [61]. Zegura et al. [61] introduce a comprehensive model that includes several previous models.

5 Internet Topology Generating Models and Tools

An abstraction or model of the actual Internet topology is important to the understanding of how the topology is formed and how Internet topology is going to be like in the future. Generated topologies are useful in assessing proposed solutions (such as routing protocols) and provide rigous foundations of analyses of how the results scale or how they might change with a different topology.

5.1 Early models

The most simple model is probably the **Pure Random** model. In this model, a set of nodes is distributed in a plane, and an edge is added between each pair of the nodes with a fixed probability p. Although this model does not explicitly attempt to reflect any structure of real networks, it is attractive for its simplicity.

The **Waxman** [59] model, on the other hand, adds edges with a probability that is some function of the distance between the node. This probability for an edge between node u and node v is given by:

$$P(u,v) = \alpha e^{-d/(\beta L)} \tag{1}$$

where $0 < \alpha$, $\beta \leq 1$, d is the Euclidean distance from u to v, and L is the maximum distance between any two nodes. There are several variations in of the Waxman model[62][63][61].

The **Transit-Stub** [64] method tries to impose a more Internet-oriented hierarchical structure as follows. A connected random graph is first generated (e.g. using the Waxman method described above). Each node in that graph represents an entire Transit domain. Each Transit domain node is expanded to form another connected random graph, representing the backbone topology of that transit domain. Next, for each node in each transit domain, a number of connected random graphs are generated, representing Stub domains that are attached to that transit node. Finally, some extra connectivity is added, in the form of "back-door" links between pairs of nodes, where a pair consists of a node from a transit domain and another from a stub domain, or one node from each of two different stub domains. By having nodes of different types, it is possible to generate large sparsely-connected Internet-like topologies with typically low node degrees. **GT-ITM**(Georgia Tech Internetwork Topology Models) is a tool to generate the Transit-Stub networks.

The problem of these early models is that they do not follow power-laws as shown in the real Internet instances. Medina et. al. [40] tested the generated topologies from Waxman and Transit-Stub, and they found both of them exhibit the absence or weak presence of the power-law.

5.2 Pure power-law models

Since the discovery of power-laws by Faloutsos brothers[1], the main focus of generating an Internet-like topology has shifted to matching the exhibited power-law in the Internet.

Palmer et. al. [65] proposed the **PLOD** (Power Law Out-Degree) model. In this model, a degree credit is first assigned to each node in a graph with a given number of nodes. The degree distribution complies with the appropriate power-laws. Then an edge placement loop is executed: it randomly picks two nodes and assigns an edge if they are not connected and each node still has remaining degree credit. After an edge is assigned, the degree credit of the nodes incident to the edge will be deducted accordingly. The loop continues until there are no more pairs of nodes that fulfill the condition.

The concept of **PLRG** (Power Law Random Graph) [66] was proposed by Aiello et al. in the year 2000, and therefore this model is also sometimes called **Model A**. In this model, a random graph is produced with power law degree distribution depending on two parameters, which roughly delineate the size and density but they are natural and convenient for describing a power law degree sequence. The power law random graph model $P(\alpha, \beta)$ is described as follows. Let y be the number of nodes with degree x. $P(\alpha, \beta)$ assigns uniform probability to all graphs with $y = e^{\alpha}/x^{\beta}$, where α is the intercept and β is the (negative) slope when the degree sequence is plotted on a log-log scale. After the degree distribution is defined, a set, L, which contains $deg(v)$ distinct copies of each node v, will be formed. Then a random matching of the elements of L is chosen. For two nodes u and v, the number of edges joining u and v is equal to the number of edges in the matching of L joining copies of u to copies of v. The graph formed in the end is the PLRG.

These power-law "matchers" do not attempt to answer how a graph comes to have a power law degree sequence. Rather, they take that as a given. Surprisingly, these method seem to be able to match many other topology properties [67] of the real Internet.

5.3 Dynamic growth models

In contrast to the pure power-law model, the dynamic growth models try to generate the Internet topology graph by simulating the growth of the Internet.

Barabasi and Albert [68] proposed a generic model (or **BA Model**) for network growth:

1. **Incremental growth**: The network expands continuously by the addition of new nodes.

2. **Preferential attachment**: A new node attaches preferentially to nodes that are already well connected.

In more detail, the network begins with a small number ($m0$) of connected nodes. New nodes are added to the network one at a time. The probability $p(v)$ that a new node is connected to an existing node v is determined as the following:

$$p(v) = d_v / \sum_j d_j \qquad (2)$$

where d_v is the degree of node v and $\sum_j d_j$ is the sum of degrees of all existing nodes. In BA model, heavily linked nodes tend to quickly accumulate even more links, while nodes with only a few links are unlikely to be chosen as the destination for a new link. The new nodes have a "preference" to attach themselves to the already heavily linked nodes. This is so called "rich-get-richer" phenomenon.

The **AB** model [69] extents the BA model by adding a third rewiring operation called "rewiring". The rewiring operation consists of choosing m links randomly and re-wire each end of them according to the same preference rule used in the BA model.

Bu et. al. [43] found that the graphs generated by PLRG, BA and AB models have different characteristic values real Internet graph in terms of path length and clustering coefficient. They proposed **GLP** (Generalized Linear Preference) [43], in which the probability p is:

$$p(v) = (d_v - \beta) / \sum_j (d_j - \beta) \qquad (3)$$

where $beta \in (-\infty, 0)$ is a tunable parameter. The smaller the value of β is, the less preference gives to high degree nodes.

All these dynamic growth models produces graphs with power-law distribution. However, it is still hard for these models to capture every topological property of the Internet. Authors in [67] show that even GLP does not follow some hierarchical properties of the Internet.

5.4 Sampling

All aforementioned models attempt to grow a graph, an approach called "constructive". One weakness of these constructive methods lies in their dependence on the principles of construction, and the choice of parameter values. Furthermore, they often focus on matching a certain number of topology properties, while fail to match some other. To address these problems, Krishnamurthy et al. [70] took a "reductive" approach: instead of trying to construct a graph, they try to "sample" real topologies to produce a smaller graph. The idea is that hopefully the original properties, either well-know or unnoticed, can be kept during the process of reduction. In more detail, they propose several reduction methods:

DRV (Deletion of Random Vertex): Remove a random vertex, each with the same probability. **DRE** (Deletion of Random Edge): Remove a random edge, each with the same probability. **DRVE** (Deletion of Random Vertex or Edge): Select a vertex uniformly at random, and then delete an edge chosen uniformly at random from the edges incident on this vertex. **DHYB-w** (Hybrid of DRVE/DRE): In this method, DRVE is executed with probability w and

DRE is executed with probability $1 - w$, where $w \in [0, 1]$. This method was motivated by the study showing that sometimes DRVE and DRE had opposite performances with respect to different metrics.

The topologies sampled by both DRV and DRE methods are mathematically proved to follow power-law degree distribution. By comparing experimental data, the authors in [70] concluded that DHYB-0.8 is the best reduction method, and it also compares favorably to graph generation methods proposed previously in the literature. These sampling methods are successful to reduce a topology down to 30% of the original size. Beyond that the statistical confidence is found low.

5.5 Topology Generation Tools

BRITE (Boston university Representative Internet Topology gEnerator) [71]is a universal topology generator. It implements a single generation model that has several degrees of freedom with respect to how the nodes are placed in the plane and the properties of the interconnection method to be used. With difference parameter settings, BRITE can generate either Waxman model or BA mode.

Inet[72] is an AS level Internet topology generator. Inet aims at reproducing the connectivity properties of Internet topologies as power-laws and with other improvements. It initially assigns node degrees from a power-law distribution and then proceed to interconnect them using different rules. The current version Inet-3.0 improves from their previous versions by creating topologies with more accurate degree distributions and minimum vertex covers as compared to Internet topologies. Inet-3.0's topologies still do not well represent the Internet in terms of maximum clique size and clustering coefficient. These related problems stress a need for a better understanding of Internet connectivity and will be addressed in the future work.

6 Conceptual Models for the Internet Topology

The Internet topology is large, complex and constantly changing. Even with the introduction of power-laws, which appears as a necessary though not sufficient condition for a topology to be realistic, a conceptual model of the topology [22][2][38] is still hard to get. Although the Internet is widely believed to be hierarchical by construction, it is too interconnected for an obvious hierarchy[45]. Several efforts to visualize the router-level topology have been made [73][8], however they can not be recreated manually and they do not provide a memorable model.

One goal here is to develop an effective conceptual model: a model that can be easily drawn by hand, while at the same time, it captures significant macroscopic properties. The Jellyfish [67] and Medusa [74] models are two conceptual models proposed for the inter-domain Internet topology.

	Instance					
	Int-11-1997		Int-06-2000		Int-07-2003	
Layer No	Nodes	% of Nodes	Nodes	% of Nodes	Nodes	% of Nodes
Core/Layer-0	8	0.23	14	0.176	13	0.08
Layer-1	1354	44.90	3659	46.25	7330	46.27
Layer-2	1202	39.866	3090	39.05	7116	45.51
Layer-3	396	13.134	1052	13.29	1078	6.89
Layer-4	43	1.425	86	10.87	96	0.61
Layer-5	12	0.398	10	0.12	1	0.0063

Table 3: Distribution of nodes in layers for three Internet instances.

6.1 The Jellyfish model

The Jellyfish model classifies ASes into different hierarchical layers. The highest layer is called the **Core**, which can be constructed in the following procedure. First sort all ASes in non-increasing degree order. The highest degree node is selected as the first member of the Core. Then each of the rest ASes is examined in that order; a node is added to the Core only if it forms a clique with the nodes already in the Core. In other words, the new node must connect to all the nodes already in the core. The procedure stops when no more node can be added. The constructed Core is a clique of high-degree ASes, but not necessarily the maximal clique of the graph. The Core is a starting point to construct a Jellyfish topology, and the ASes in the core are probably the most important ASes in the Internet since they have high degrees. The rest of the nodes are defined according to their proximity to the Core. The first **layer** is defined as all the ASes adjacent to the Core. Similarly, the second layer is defined as the non-labeled neighbors of the first layer. By repeating this procedure, six layers can be identified from the instances of Internet AS-level topology if the Core is counted as layer zero. Table 3 shows the number and percentage of ASes with each layer for three Internet topology instances at different time.

The Jellyfish model also studies separately the one-degree ASes. This is because, although one-degree nodes do not provide connectivities to the rest of the network, they have significant numbers. In fact, 35%-45% of the ASes in the Internet are one-degree. In the Jellyfish model, each layer is separated into two classes: a) the multiple-degree or **shell** nodes, and b) the one-degree or **hang** nodes. The one-degree nodes hanging from k-th shell are referred as the k-th hang class. For example, shell-0 is the core, and its one-degree neighbors are denoted as hang-0, while the rest of the neighbors constitute shell-1. Naturally, the number of ASes in the layers, shells and hangs have the following relationship:

$$Layer_k = Shell_k + Hang_{k-1}$$

Table 4 shows the size of each group of nodes in the classification.

The conceptual Jellyfish model is described by the layer-shell-hang classification. The Core is the center of the head of the jellyfish surrounded by shells

	Instance					
	Int-11-1997		Int-06-2000		Int-07-2003	
Layer ID	Nodes	% of Nodes	Nodes	% of Nodes	Nodes	% of Nodes
Core/Shell-0	8	0.23	14	0.176	13	0.08
Hang-0	465	15.42	798	10.08	1174	7.5
Shell-1	889	29.49	2861	36.16	6156	39.37
Hang-1	623	20.66	1266	16	2821	18.04
Shell-2	579	19.2	1824	23.05	4295	27.47
Hang-2	299	9.92	662	8.36	808	5.16
Shell-3	97	3.22	390	4.92	270	1.72
Hang-3	41	1.36	74	0.93	84	0.53
Shell-4	2	0.66	12	0.15	12	0.07
Hang-4	12	0.4	10	0.12	1	0.006

Table 4: Distribution of nodes in shell and hang classes.

of nodes. Figure 6 shows a graphical illustration of this model. The hang nodes form the tentacles of the jellyfish. The length of the tentacle represents the concentration of one-degree neighbors for each shell.

Besides being a conceptual model, the Jellyfish represents some invariant properties of the Internet topology. (1) Core: The topology has a core of highly connected important nodes, which is represented by the center of the jellyfish cap. (2) Center-heavy: Approximately 80% of the Ases are layer-1, layer-2 and layer-3. (See Table 3.) (3) Node distance: Distances between ASes are small; maximum distance less than 11 hops, and 80% of the ASes are within 5 hops. (4) Edges type: Approximately 70% of the edges are between different node layers. The rest are horizontal to the hierarchy providing connectivity between nodes of the same layer. (5) One-degree nodes: There is a non-trivial percentage (35%-45%) of one-degree nodes. (See Table 4.)

6.2 Jellyfish tells the difference

As shown in the previous section, the Internet topology fits in the Jellyfish profile. However, not every graph can be modeled as a jellyfish. For example, if a tree with N nodes were to modeled into a jellyfish, the number of shells would be proportional to $O(\log N)$. This does not fit into the Jellyfish profile of Internet, where the number of shells is constant at 5 despite the rapid growth of the number of nodes. This provides an opportunity to use the Jellyfish model as a test of the realism of Internet like graphs.

Siganos et. al. [67] use the Jellyfish model to test the GLP methodology proposed in [43], and the PLRG approach proposed in [66]. The GLP approach depends on the preferential model. On the other hand, the PLRG generator is based on an interesting theoretical model for scale free graphs, and takes the degree distribution as a given. In [43], these two generators were compared and it was concluded that the best generator was the GLP. PLRG was shown to

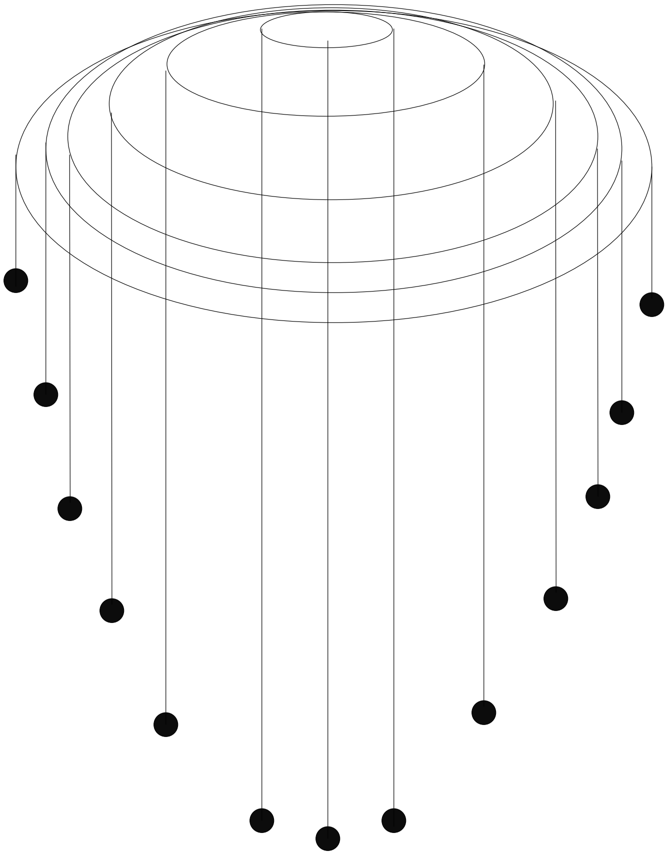

Figure 6: The Internet topology as a jellyfish.

fail capturing properties like the characteristic path length and the clustering coefficient. However, by using Jellyfish model, Siganos et. al. [67] were able to show that GLP does not capture the macro structure by using jellyfish. Incidentally, PLRG seems to pass the test, although it fails other properties.

In more detail, two graphs are generated by the GLP model and the PLRG model respectively. Both of them have similar number of nodes to an Internet topology instance in Jun 2000. In table 5, these graphs are decomposed using the jellyfish model. These results clearly show that the generated graph using the GLP methodology is qualitatively different than the Internet graph. First, the Core of the network is much bigger in GLP (21) compared to the Internet (14). Second, The number of hanging nodes (degree one) in GLP far out-exceeds the number of shell nodes. The ratio is approximately 70% hanging nodes to 30% shell nodes. In the case of the Internet, this ratio is the opposite. Third, the GLP topology has only up to 5 layers, with the 5th layer having only 3 members, while the real Internet has 6 layers. On the other hand PLRG seems to maintain

	Instance					
	GLP		Int-06-2000		PLRG	
Layer ID	Nodes	% of Nodes	Nodes	% of Nodes	Nodes	% of Nodes
Core/Shell-0	21	0.2	14	0.176	11	0.13
Hang-0	1885	23.82	798	10.08	565	7.1
Shell-1	1672	21.13	2861	36.16	2346	29.6
Hang-1	3371	42.6	1266	16	1298	16.4
Shell-2	688	8.7	1824	23.05	2305	29.13
Hang-2	221	2.79	662	8.36	525	6.6
Shell-3	3	0.037	390	4.92	325	4.1
Hang-3	3	0.037	74	0.93	125	1.5
Shell-4	0	0	12	0.15	41	0.51
Hang-4	0	0	10	0.12	23	0.29

Table 5: Distribution of nodes in shell and hang classes.

similar structure according to the jellyfish model. The only differences between PLRG and the Internet is that the clique is smaller, having only 11 nodes, and that there is a slightly smaller shell-1 and a bigger shell-2.

6.3 The Medusa model

One problem of the Jellyfish model is that the identities of the Jellyfish Core are not particularly robust when the completeness of the Internet topology is uncertain. Carmi et al. [74] found that by adding or deleting an edge, the ASes in the Jellyfish Core could change up to 25%, mostly affecting some European ASes. To address the problem, they proposed a model called *Medusa*. The Medusa model depends on an informative functional decomposition of the Internet AS called k-pruning, which proceeds as follows:

First, each AS with only one neighbor is removed. The link to that neighbor along with the node is removed as well. As this pruning proceeds, further nodes with one neighbor (or fewer) may be exposed. They will be removed until there is no 1-degree ASes in the remaining graph. ASes removed in this way make up what is called *1-shell*. The remaining graph is called *2-core*. Second, the pruning process is repeated and it is characterized by an index k. For example, when $k = 2$, all nodes with 2 neighbors will be removed from the 2-core, and the removed nodes in this step is called *2-shell*. The process continues, eliminating any nodes reduced to a degree of 2 (or fewer) by this pruning, until all nodes remaining have 3 or more neighbors. The remaining graph is called *3-core*. The process is repeated to identify the 3-shell and 4-core, and so on. The process stops at the point when no further nodes remain. The last nonempty k-core provides a very robust and natural definition of the heart or nucleus of any communications network. Last, the *k-crust* is defined as the union of the nodes in the 1 through k shells, and the links that join them. The $k - 1$ crust is the complement of the k-core.

For small k, the crusts consist of many small clusters of connected sites. For sufficiently large k, the largest connected cluster of a k-crust consists of a significant fraction of the whole k-crust, while no smaller cluster contains more than a few nodes. The change occurs at a well-defined threshold value of k. There is a significant fraction of the nodes within each large-k crust which is not part of its largest cluster, and remain isolated. Thus the AS graph (or any similar scale-free network) can be decomposed into three distinct components:

1. The nucleus (the innermost k-core)
2. The giant connected component of the last crust, in which only the nucleus is left out
3. The isolated components of the last crust, nodes forming many small clusters. These connect to the connected component of the last crust only through the nucleus

These three classes of nodes are quite different in their functional role within the Internet. The nucleus plays a critical role in BGP routing, since its nodes lie on a large fraction of the paths that connect different ASes. It allows redundancy in path construction, which gives immunity to multiple points of failure. The connected component of the large-k crusts could be an effective substrate on which to develop additional routing capacity, for messages that do not need to circle the globe. Finally, the isolated nodes and isolated groups of nodes in the last crust essentially leave all routing up to the nodes in the nucleus of the network. Because all their message traffic passes through the nucleus, even when the destination is relatively close by, they may be contributing unnecessary load to the most heavily used portions of the Internet. The relative size of this component could be a key indicator of the evolution of the topography of the Internet.

This model can be visualized as Figure 7. The core of the Medusa includes the most important nodes that are found in the core and the first ring of the Jellyfish's mantle. The Jellyfish has relatively few rings around its core, while the Medusa's mantle is more extended and differentiated. The tendrils hanging from the Jellyfish (leaf nodes) descend mostly from the core, but also from all the other rings, while all the tendrils of the Medusa are, by construction, attached to its nucleus.

7 The Complete Internet Topology

An accurate topology model would be important for simulating, analyzing, and designing the future protocols effectively [4]. With an accurate Internet AS-level topology, first, one can design and analyze new inter-domain routing protocols, such as HLP [75], that take advantage of the properties of the Internet AS-level topology. Second, one can create more accurate models for simulation purposes [76]. Third, one can analyze phenomena such as the spread of viruses

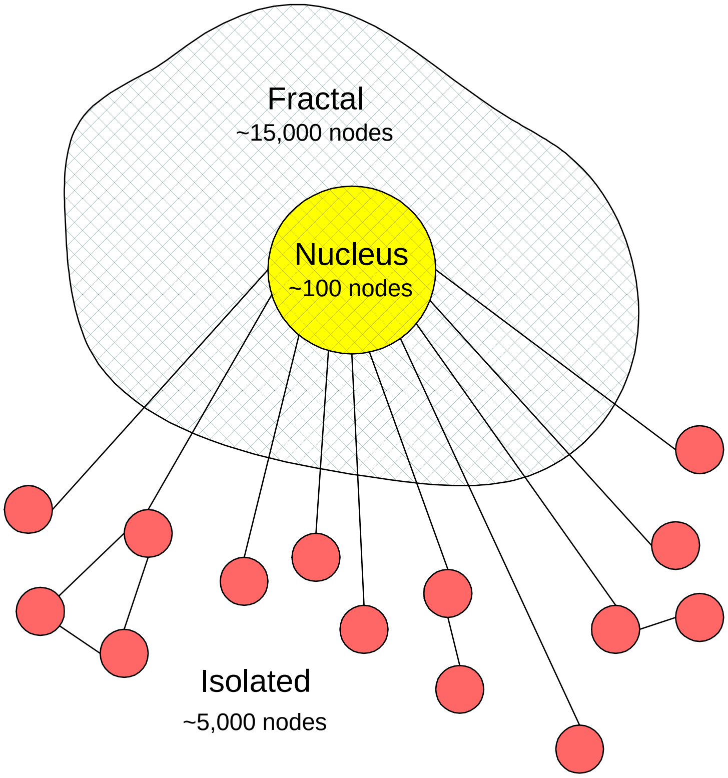

Figure 7: The Internet topology as a medusa.

[77][78], more accurately. In addition, the current initiatives of rethinking and redesigning the Internet and its operation from scratch would also benefit from such an accurate Internet topology.

7.1 Toward finding a representative snapshot of Internet topology

Developing an accurate representation of the Internet topology at the AS level remains as a challenge despite the recent flurry of studies [52][18][11][19][79][80][17][16]. One of the major problems is that, although ASes are generally presented completely in most of the Internet topology sources, the edges among the ASes are not. As seen in the previous section, each of topological sources has its own advantages, but each of them also provides an incomplete, sometimes inaccurate view of the Internet AS topology; these views are often complementary.

Recently, He et. al [53] present a systematic framework for extracting and synthesizing the AS level topology information from different sources. Instead of simply taking the union of all resources, a careful synthesis and cross-validation is performed. In additional to the sources mentioned above, they also utilize information gathered from IXPs (Internet Exchange Points), which have not received attention in terms of Internet topology discovery, although they play a major role in the Internet connectivity.

They identify and validate several properties of the missing AS links. (1)most of the missing AS edges are of the peer-to-peer type, (2) many of the missing AS edges from BGP tables appear in IRR, and (3) most of the missing peer-to-peer AS edges are incident at IXPs.

Their work consists four steps.

First, BGP routing tables are compared. They consider the AS edges derived from multiple BGP routing table dumps[18], and compare them to the Routeview data (OBD). The question to answer is what is the information that the new BGP tables bring. Table 6 lists a portion of the collection of BGP table dumps that were collected in May 2005. One observation here is that, about 80% of the missing links that do not appear in a single table dump (OBD) but appear in a collection of table dumps (BD) are peer-to-peer type. For example, among 8702 edges in BD but not in OBD, 7183 of them are classified as peer-to-peer type.

Second, they systematically analyze the IRR data and identify topological information that seems trustworthy by Nemecis[21]. They follow a conservative approach, given that IRR may contain some out-dated and/or erroneous information. They do not accept new edges from IRR, even after our first processing, unless they are confirmed by traceroutes by using public traceroute servers. Overall, they find that IRR is a good source of missing links. For example, they discover that more than 80% of the new edges found in the extra tables already exist in IRR [20]. On the other hand, IRR has still have significantly more edges.

Third, they identify the ASes which participate at Internet Exchange Points (IXPs). An IXP is a relative low cost solution for an AS to peer with many other

Table 6: A collection of BGP table dumps

Route collector or Router server name	# of Nodes	# of Edges	# of edges with type inferred			edges not in OBD	edges not in OBD w/ type		
			total	p-p	p-c		total	p-p	p-c
route-views(*OBD*)	19843	42643	42570	5551	36766	0	0	0	0
route-views2	19837	41274	41230	4464	36514	1029	1028	835	191
route-views.eqix	19650	34889	34876	1027	33640	674	674	530	143
route-views.linx	19655	37259	37246	3246	33765	2511	2511	2188	319
route-views.isc	19753	36152	36139	1915	34004	784	783	663	118
rrc00.ripe	19770	36479	36465	1641	34605	655	654	543	111
rrc01.ripe	19640	34193	34180	1121	32855	617	617	512	105
rrc03.ripe	19737	39147	39129	3850	35042	3233	3228	2609	616
rrc05.ripe	19765	32676	32659	1122	31324	1095	1091	658	432
rrc07.ripe	19618	32811	31797	1219	30394	804	803	724	79
rrc12.ripe	19628	33841	33827	2024	31606	1611	1610	1417	193
Total(*BD*)	19950	51345	51259	12734	38265	8702	8689	7183	1499

peers who are also participants at the same IXP. The exhaustive identification of IXP participants has received limited attention. Most previous work focuses on identifying the existence of IXPs. The finding here is that many of the ASes incident to the peer-to-peer edges missing from the different data sets are IXP participants. Note that even if two ASes peer at the same IXP, that does not necessarily mean there is an AS edge between these two ASes, because this totally depends on peering agreement between these two ASes. Therefore, in order to test whether or not these missing edges are indeed at the IXPs, they proceed to the next step.

Fourth, they use their traceroute tool, RETRO, to verify potential edges from IRR and IXPs. RETRO is a tool that collects public traceroute server configurations, send out traceroute requests, and collect traceroute results dynamically. They confirm the existence of many potential edges identified in the previous steps. The results show that more than 94% of the RETRO-verified AS edges in IRR indeed go through IXPs. They even discover edges that were not previously seen in either the BGP table dumps or IRR. In total, 300% more peer-to-peer links than those in the conventional BGP table dumps from Routeviews have been validated.

7.2 Towards finding Internet backup links

One limitation of the previous method is that it ultimately depends on traceroute to verify the existence of a suspected edge. It is plausible if the suspected edge is a primary link, which means they exist most of the time. If an suspected edge is a backup, and it does not show unless some other links break down, it is unlikely to be witnessed by traceroute. Even one can keep probing for a long period of time, it is hard to tell if a link observed before is still existing currently.

Recently, active BGP probing [17] has been proposed for identifying backup AS links. The main idea is to inject false AS path loops for an unused IP block. Since AS path loops are prohibited in inter-domain routing, BGP routers are forced to switch to backup links for this unused IP block. These links can be observed from any route collectors, such as Routeviews or RIPE/RIS. This probing technique does not affect normal Internet routing because every change is restricted in the unused IP block.

In more detail, the principle of active BGP probing is the following. An active probing AS announces one of its prefixes with AS-paths including a number of other ASes. These ASes, due to loop detection, will not use or propagate the announcement. Then if there is any alternative path available, it will show up. To avoid influencing AS-path length, the prohibited ASes are placed in an AS-set at the end of the path. For example, to stop its announcement from being propagated by ASes 1, 2, and 3, an AS (say AS12654) might announce one of its prefixes with an AS-path of [12654 {1,2,3}]. This allows AS 12654 to discover who propagates its announcements, find backup paths, and deduce the policies of other ASes with respect to its prefixes. By proper constructing the "prohibited" AS sets, one may be able to discover all backup links visible to the probing AS.

Note, due to BGP export policies and very limited number of probing ASes, this method majorly discovers provider-customer type backup links. This method and the method of discovering missing peer-to-peer links [53] are complementary to each other.

8 Future Directions

degree-degree correlation sigcomm 2006, topology with directions (Gao) Topology generator and sampler for BGP simulation. Evolution (infocom 2006) Router-level topology discovery (NetDimes, iplane, rocketfuel)

9 Bibliography

Primary Literature(cited)

References

[1] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the Internet topology. *ACM SIGCOMM*, pages 251–262, Sep 1-3, Cambridge MA, 1999.

[2] J.-J. Pansiot and D Grad. On routes and multicast trees in the Internet. *ACM SIGCOMM Computer Communication Review*, 28(1):41–50, January 1998.

[3] R. Govindan and A. Reddy. An analysis of Internet Inter-domain topology and route stability. *Proc. IEEE INFOCOM*, Kobe, Japan, April 7-11 1997.

[4] S. Floyd and V. Paxson. Difficulties in simulating the Internet. *IEEE Transaction on Networking*, Aug 2001.

[5] Oregon routeview project, http://www.routeviews.org.

[6] Ripe route information service, http://www.ripe.net/ris.

[7] V. Jacobson. Traceroute. Internet Measurement Tool, 1995.

[8] Skitter, http://www.caida.org/tools/measurement/skitter/.

[9] The cooperative association for internet data analysis, http://www.caida.org.

[10] N. Spring, R. Mahajan, D. Wetherall, and T. Anderson. Measuring ISP topologies with rocketfuel. *IEEE/ACM Trans. Netw.*, 12(1):2–16, 2004.

[11] Y. Shavitt and E. Shir. DIMES: Let the Internet Measure Itself. *ACM SIGCOMM Computer Communication Review (CCR)*, October 2005.

[12] Z. Mao, J. Rexford, J. Wang, and R. Katz. Towards an accurate AS-level traceroute tool. In *Sigcomm*, 2003.

[13] Z. Mao, D. Johnson, J. Rexford, J. Wang, and R. Katz. Scalable and accurate identification of AS-Level forwarding paths. In *Infocom*, 2004.

[14] M. Crovella A. Lakhina, J. W. Byers and I. Matta. Sampling biases in ip topology measurements. In *IEEE Infocom*, 2003.

[15] D. Achlioptas, A. Clauset, D. Kempe, and C. Moore. On the bias of traceroute sampling, or power-law degree distributions in regular graphs. In *STOC*, 2005.

[16] R. Cohen and D. Raz. The Internet Dark Matter – on the Missing Links in the AS Connectivity Map. In *IEEE Infocom*, 2006.

[17] L. Colitti, G. Di Battista, M. Patrignani, M. Pissonia, and M. Rimondini. Investigating prefix propagation through active BGP probing. In *IEEE ISCC*, 2006.

[18] B. Zhang, R. Liu, D. Massey, and L. Zhang. Collecting the Internet AS-level Topology. *ACM SIGCOMM Computer Communication Review(CCR)*, January 2005.

[19] X. Dimitropoulos, D. Krioukov, and G. Riley. Revisiting Internet AS-Level Topology Discovery. In *PAM*, 2005.

[20] Internet routing registry, http://www.irr.net.

[21] G. Siganos and M. Faloutsos. Analyzing BGP Policies: Methodology and Tool. In *IEEE Infocom*, 2004.

[22] G. Siganos, M. Faloutsos, P. Faloutsos, and C. Faloutsos. Power-laws of the Internet topology. *IEEE/ACM Trans. on Networking*, 1(4):514–524, 2003.

[23] V. Pareto. Cours d'economic politique. *Dronz,Geneva Switzerland*, 1896.

[24] G.K. Zipf. *Human Behavior and Principle of Least Effort: An Introduction to Human Ecology.* Addison Wesley, Cambridge, Massachusetts, 1949.

[25] W.E. Leland, M.S. Taqqu, W. Willinger, and D.V. Wilson. On the self-similar nature of ethernet traffic. *IEEE Transactions on Networking*, 2(1):1–15, February 1994. (earlier version in SIGCOMM '93, pp 183-193).

[26] V. Paxson and S. Floyd. Wide-area traffic: The failure of poisson modeling. *IEEE/ACM Transactions on Networking*, 3(3):226–244, June 1995. (earlier version in SIGCOMM'94, pp. 257-268).

[27] M. Crovella and A. Bestavros. Self-similarity in World Wide Web traffic, evidence and possible causes. *SIGMETRICS*, pages 160–169, 1996.

[28] R. Albert, H.Jeong, and A.L. Barabasi. Diameter of the world wide web. *Nature*, 401, 1999.

[29] R. Kumar, P. Raghavan, S. Rajagopalan, D. Sivakumar, A.Tomkins, and E. Upfal. The web as a graph. *ACM Symposium on Principles of Database Systems*, 2000.

[30] M.Jovanovic. Modeling large-scale peer-to-peer networks and a case study of gnutella. *Master thesis,University of Cincinnati*, 2001.

[31] J. Chuang and M. Sirbu. Pricing multicast communications: A cost based approach. *In Proc. of the INET'98*, 1998.

[32] G. Philips, S. Shenker, and H. Tangmunarunkit. Scaling of multicast trees: Comments on the chuang-sirbu scaling law. *ACM SIGCOMM*, Sep 1999.

[33] T. Wong and R. Katz. An analysis of multicast forwarding state scalability. *International Conference on Network Protocols*, 2000.

[34] P. Van Mieghem, G. Hooghiemstra, and R. van der Hofstad. On the efficiency of multicast. *IEEE/ACM Transactions on Networking*, 9, 2001.

[35] Citeseer, http://citeseer.ist.psu.edu/articles1999.html.

[36] S. Jamin, C. Jin, Y. Jin, D. Raz, Y. Shavitt, and L. Zhang. On the placement of Internet instrumentation. *Proc. IEEE INFOCOM*, Tel Aviv, Israel, March 2000.

[37] R. Govindan and H. Tangmunarunkit. Heuristics for Internet map discovery. *Proc. IEEE INFOCOM*, Tel Aviv, Israel, March 2000.

[38] Damien Magoni and Jean Jacques Pansiot. Analysis of the autonomous system network topology. *ACM Computer Communication Review*, July 2001.

[39] A. Medina, I. Matta, and J. Byers. On the origin of powerlaws in Internet topologies. *CCR*, 30(2):18–34, April 2000.

[40] A. Medina, A. Lakhina, I. Matta, and J. Byers. Brite:an approach to universal topology generation. *MASCOTS*, 2001.

[41] Christopher R. Palmer and J. Gregory Steffan. Generating network topologies that obey power laws. *IEEE Globecom*, 2000.

[42] Cheng Jin, Qian Chen, and Sugih Jamin. Inet: Internet topology generator. *Techical Report UM CSE-TR-433-00*, 2000.

[43] T.Bu and D. Towsley. On distinguishing between Internet power law topology generators. *Infocom*, 2002.

[44] S.H. Yook, H.Jeong, and A. Barabasi. Modeling the Internet's large-scale topology. *PNAS*, 99(21), 2002.

[45] H. Tangmunarunkit, R. Govindan, S. Jamin, S. Shenker, and W. Willinger. Network Topology Generators: Degree based vs. Structural. In *ACM Sigcomm*, 2002.

[46] Q. Chen, H. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger. The Origin of Power Laws in Internet Topologies Revisited. In *Infocom*, 2002.

[47] S. Jaiswal, A. Rosenberg, and D. Towsley. Comparing the structure of power law graphs and the Internet AS graph. In *ICNP*, 2004.

[48] H. Chou. A note on power-laws of Internet topology.

[49] L.A.Adamic. Zipf, power-laws, and pareto - a ranking tutorial. *http://www.parc.xerox.com/iea/*, 2000.

[50] D. M. Cvetkoviċ, M. Boob, and H. Sachs. *Spectra of Graphs*. Academic press, 1979.

[51] M.Mihail and C.H.Papadimitriou. On the eigenvalue power law. *Random*, 2002.

[52] H. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger. Towards capturing representative AS-level Internet topologies. *Computer Networks*, 44(6):737–755, 2004.

[53] Y. He, G. Siganos, M. Faloutsos, and S. Krishnamurthy. A Systematic Framework for Unearthing the Missing Links: Measurements and Impact. In *USENIX NSDI*, 2007.

[54] L. Gao. On inferring autonomous system relationships in the Internet. In *IEEE Global Internet*, November 2000.

[55] J. Xia and L. Gao. On the evaluation of as relationship inferences. In *IEEE Globecom*, November 2004.

[56] G.D. Battista, M. Patrignani, and M. Pizzonia. Computing the types of the relationships between Autonomous Systems, 2003.

[57] X. Dimitropoulos, D. Krioukov, M. Fomenkov, B. Huffaker, Y. Hyun, kc claffy, and G. Riley. As relationships:inference and validation. *ACM SIGCOMM Computer Communication Review (CCR)*, January 2007.

[58] F. Wang and L. Gao. Inferring and characterizing internet routing policies. In *ACM IMW*, 2003.

[59] B. M. Waxman. Routing of multipoint connections. *IEEE Journal of Selected Areas in Communications*, pages 1617–1622, 1988.

[60] M. Doar. A better model for generating test networks. *Proc. Global Internet, IEEE*, Nov. 1996.

[61] E. W. Zegura, K. L. Calvert, and M. J. Donahoo. A quantitative comparison of graph-based models for internetworks. *Transactions on Networking*, 5(6):770–783, December 1997.

[62] M. Doar and I. Leslie. How bad is naive multicast routing? *Proc. IEEE INFOCOM*, pages 82–89, 1993.

[63] L. Wei and D. Estrin. The trade-offs of multicast trees and algorithms. *International Conference on Computer Communications and Networks*, 1994.

[64] Ken Calvert, Matt Doar, and Ellen W. Zegura. Modeling internet topology. *IEEE Communication Magazine*, June 1997.

[65] C. R. Palmer and J. G. Stefan. Generating network topologies that obey powerlaws. *Proceedings of the Global Internet Symposium, GLOBECOM 2000*, 2000.

[66] W. Aiello, F. Chung, and L. Lu. A random graph model for massive graphs. *STOC*, 2000.

[67] G. Siganos, S. Tauro, and M. Faloutsos. Jellyfish: A conceptual model for the as internet topology. *Journal of Communications and Networks*, 8(3):339–350, 2006.

[68] A. Barabasi and R. Albert. Emergence of scaling in random networks. *Science*, 8, October 1999.

[69] R. Albert and A. Barabasi. Topology of complex networks:local events and universality. *Phys.Review*, 85, 2000.

[70] V. Krishnamurthy, M. Faloutsos, M. Chrobak, L. Lao, J-H. Cui, and A.G. Percus. Reducing large internet topologies for faster simulations. In *Networking*, 2005.

[71] Brite, http://www.cs.bu.edu/brite/.

[72] Inet topology generator, http://topology.eecs.umich.edu/inet/.

[73] Bill Cheswick and Hal Burch. Internet mapping project. Wired Magazine, December 1998. See http://cm.bell-labs.com/cm/cs/who/ches/map/index.html.

[74] Shai Carmi, Shlomo Havlin, Scott Kirkpatrick, Yuval Shavitt, and Eran Shir. Medusa - new model of internet topology using k-shell decomposition. 2006.

[75] L. Subramanian, M. Caesar, C. T. Ee, M. Handley, M. Mao, S. Shenker, and I. Stoica. HLP: A Next-generation Interdomain Routing Protocol. In *ACM Sigcomm*, 2005.

[76] O. Maennel and A. Feldmann. Realistic BGP Traffic for Test Labs. In *ACM Sigcomm*, 2002.

[77] K. Park and H. Lee. On the effectiveness of route-based packet filtering for distributed DoS attack prevention in power-law Internets. In *ACM Sigcomm*, Aug 2001.

[78] A. Ganesh, L. Massoulie, and D. Towsley. The Effect of Network Topology on the Spread of Epidemics. In *IEEE infocom*, 2005.

[79] P. Mahadevan, D. Krioukov, M. Fomenkov, B. Huffaker, X. Dimitropoulos, kc claffy, and A. Vahdat. The Internet AS-Level Topology: Three Data Sources and One Definitive Metric. *ACM SIGCOMM Computer Communication Review (CCR)*, January 2006.

[80] H. Chang, S. Jamin, and W. Willinger. To Peer or not to Peer: Modeling the Evolution of the Internet's AS Topology. In *IEEE Infocom*, 2006.

Books and Reviews(uncited)